EXHIBIT 44
[FILED UNDER SEAL]

THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
SHERMAN DIVISION

STATE OF TEXAS, ET AL., **Plaintiffs,** **v.** **GOOGLE LLC,** **Defendant.**	Civil Action No.: 4:20-cv-00957 (SDJ)

EXPERT REPORT OF PAUL R. MILGROM

DATE: July 30, 2024

HIGHLY CONFIDENTIAL
SUBJECT TO PROTECTIVE ORDER

TABLE OF CONTENTS

I. ASSIGNMENT AND QUALIFICATIONS

A. Qualifications

1. My name is Paul R. Milgrom. I am the Shirley and Leonard Ely Professor of Humanities and Sciences in the Department of Economics at Stanford University and professor, by courtesy, at both the Department of Management Science and Engineering and the Graduate School of Business. I am also the chairman and co-founder of Auctionomics, which designs and assists bidders in high-stakes auctions.

2. In 2020, I was the co-recipient, with Professor Robert Wilson, of the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel, commonly known as the Nobel Prize in Economics, "for improvements to auction theory and inventions of new auction formats." As explained by the Royal Swedish Academy of Sciences: "The new auction formats are a beautiful example of how basic research can subsequently generate inventions that benefit society. The unusual feature of this example is that the same people developed the theory and the practical applications. The Laureates' ground-breaking research about auctions has thus been of great benefit, for buyers, sellers and society as a whole."[1]

3. Earlier in the same year, I was also named a Distinguished Fellow of the American Economic Association. The Distinguished Fellow citation describes me as "the world's leading auction designer, having helped design many of the auctions for radio spectrum conducted around the world in the last thirty years, including those conducted by the US Federal Communications Commission (ranging from the original simultaneous multiple

[1] The Royal Swedish Academy of Sciences, "The quest for the perfect auction," Nobelprize.org (2020), https://www.nobelprize.org/uploads/2020/09/popular-economicsciencesprize2020.pdf, at 6-7.

round auction with activity rules, to the recent incentive auction for repurposing broadcast spectrum for modern uses). His applied work in auction design and consulting has established new ways for economists to interact with the wider world. He is also a theorist of extraordinary breadth, who has provided (and still continues to provide) foundational insights not only into the theory of auctions (including his 1982 paper with Weber), but across the range of modern microeconomic theory."[2]

4. Continuing, the citation notes that "[h]is work has been widely recognized. He is a member of the National Academy of Sciences and the American Academy of Arts and Sciences. He has received major prizes, including the 2008 Nemmers Prize, the 2012 BBVA Foundation Frontiers of Knowledge Award, the 2014 Golden Goose Award (with McAfee and Wilson), the 2018 CME Group-MSRI Prize in Innovative Quantitative Applications, and the 2018 John J. Carty Award for the Advancement of Science (with Kreps and Wilson). He is the dissertation advisor of many successful economists."

5. I have been Professor at Stanford University since 1987. My prior academic appointments were at Yale University (from 1982 to 1987) and Northwestern University (from 1979 to 1983). In 2023, I was also a Distinguished Research Professor at the Simons Laufer Mathematical Sciences Institute (supported by the Alfred P. Sloan Foundation). I hold a Ph.D. in Business from Stanford University (conferred in 1979), a M.S. in Statistics from Stanford University (conferred in 1978), and an A.B. in Mathematics with high honors from University of Michigan (conferred in 1970).

[2] American Economic Association, "Paul Milgrom, Distinguished Fellow 2020" (2020), https://www.aeaweb.org/about-aea/honors-awards/distinguished-fellows/paul-milgrom.

6. At Stanford University, I teach undergraduate and graduate courses in microeconomic theory and market design. *Market design* is a field of research in economics, management science and computer science, which includes the study of auctions.[3]

7. My academic research on auctions and economic theory has been published in a number of peer-reviewed journals in economics, including *Econometrica, American Economic Review, Journal of Political Economy, Quarterly Journal of Economics, Journal of Financial Economics, Games and Economic Behavior, Journal of Economic Perspectives, Journal of Economic Theory,* and *Journal of Mathematical Economics.*

8. In the online display advertising industry, I was engaged to give commercial advice to two companies. From 2007 to 2008, I advised Yahoo! Inc., which was a leading online publisher and operator of an ad network. From 2009 to 2017, I occasionally advised OpenX, a supply-side platform, on auction design-related issues. As part of this work, I

---

[3] The National Bureau of Economics Research (NBER) has a market design working group, which it describes as one that "studies market institutions such as auctions, queues, assignment rules in school systems, clearinghouses, and tradeable permit systems. It emphasizes the role of institutional design in determining market outcomes and the well-being of market participants." NBER, Market Design (accessed Sep. 27, 2023), https://www.nber.org/programs-projects/programs-working-groups%23Groups/market-design.

The Institute for Operations Research and Management Science (INFORMS) has an "Auctions and Market Design Cluster," for which the "[a]pplications include procurement auctions, spectrum auctions, kidney exchanges, labour markets, or digital advertising markets." INFORMS, "About AMD - Auctions and Market Design" (accessed Sep. 27, 2023), http://connect.informs.org/auctionsandmarketdesign/about-us/aboutamd.

The Simons Laufer Mathematical Sciences Institute of University of California at Berkeley had a 2023 program entitled "Mathematics and Computer Science of Market and Mechanism Design." It explains that "economists and computer scientists have collaborated with mathematicians, operations research experts, and practitioners to improve the design and operations of real-world marketplaces." SLMath, "Mathematics and Computer Science of Market and Mechanism Design" (2023), https://www.slmath.org/programs/333.

At Stanford University, I teach an economics course on market design, which has coverage including "the design of platforms and exchanges, with applications to internet markets." Stanford University, Stanford Bulletin (2023), https://explorecourses.stanford.edu/search;jsessionid=zg9iqiunv63g16z0qr48bmb90?q=ECON+136%3a+Market+Design&view=catalog&filter-coursestatus-Active=on&academicYear=20232024.

co-invented an auction design that was patented by OpenX, entitled "*Impression allocation system and methods using an auction that considers losing bids*."[4]

9. In 2017-2018, I was a visiting research scholar at Google, studying the economics and pricing of cloud computing.

10. I attach my curriculum vitae in Appendix A, which provides further biographical details, including details of my previous work as an expert witness.

B. Assignment

11. The State of Texas and a group of other states (collectively, "Plaintiffs"[5]) have alleged that Google has violated the Sherman Act and state antitrust and consumer protection laws via "deceptive trade practices and anticompetitive conduct."[6]

12. I have been retained on behalf of Google LLC ("Google") to analyze and assess the economic effects of Google's online display advertising auction practices that Plaintiffs and their experts allege to be deceptive and/or anticompetitive. Specifically, I studied the economic effects of the following auction practices:

[4] Milgrom, P.R., Cunningham, S.J., & Beck, M.R. (2023). *U.S. Patent No. 11,574,358-B2.* Washington, DC: U.S. Patent and Trademark Office.

[5] The Plaintiffs consist of the States of Texas, Alaska, Arkansas, Florida, Idaho, Indiana, Louisiana, Mississippi, Missouri, Montana, Nevada, North Dakota, South Carolina, South Dakota and Utah, and the Commonwealths of Kentucky and Puerto Rico. *State of Texas et al. v. Google LLC*, Fourth Amended Complaint, May 5, 2023 ("Fourth Amended Complaint").

[6] Fourth Amended Complaint ¶ 30.

a. *Buy-side DRS*, *Project Bernanke*, *Global Bernanke* and *Alchemist*[7]: Google Ads' bid optimization programs for AdX auctions designed to maximize the value of impressions won by Google Ads, without increasing its overall revenue share.

b. *Project Bell*[8]: Google Ads' practice of reducing bids to publishers that query AdX multiple times for the same impression.

a. *Project Elmo*: a budget management feature on DV360 and Google Ads to ensure that Google made consistent bids on behalf of its advertisers across all bid requests received for a given end user within each minute.

c. *Projects Poirot and Marple*[9]: DV360's and Google Ads' programs to detect ad exchanges using non-second-price auction formats and adjust bids to maximize advertiser profits in those auctions.

d. *Dynamic Allocation (DA)*: a procedure to increase yields on publisher impressions by inserting Google's real-time auctions into publishers' earlier sequential "waterfall" allocation process.

e. *Enhanced Dynamic Allocation (EDA)*: a technology that dynamically allocates impressions between guaranteed contracts and an auction process to increase publisher revenues and improve efficiency.

[7] None of the Plaintiffs' allegations are about Buy-side DRS, but I have analyzed it in this report as it is a close precursor of Project Bernanke.

[8] In this report, I use "Project Bell" or "Bell" to refer to what Google internally called Bell v.2 (*i.e.*, not the program "Global Bernanke," which was also at times called Bell v.1).

[9] None of the Plaintiffs' allegations are about Project Marple, but I have analyzed it in this report as a closely related program to Project Poirot.

f. The so-called *"last look"*: a side effect of the way that some publishers configured header bidding to integrate with DFP causing bids from AdX bidders to be received after header bids had been received by the publisher.

g. *Reserve Price Optimization (RPO)*: a DFP feature that automatically increased floor prices for impressions when Google detected that a publisher set a floor below the revenue-maximizing level.

h. *Sell-side Dynamic Revenue Share (DRS)*: AdX's practice of varying its revenue share on individual impressions to increase the number of impressions it sold, without increasing its overall revenue share.

i. *Open Bidding*: Google's auction design to integrate real-time bids from other exchanges into the sale of impressions.

j. *Unified First Price Auction (UFPA)*: Google's auction redesign that compares bids for impressions from all bidders on the same first-price basis.

k. *Uniform Pricing Rules (UPR)*: a feature that allows publishers to configure and manage floor prices that apply equally to all buyers (*i.e.*, exchanges and other demand sources) participating in Google's UFPA.

C. Compensation

13. I am compensated at the rate of $1,800 per hour for the time I work on this matter, which is my current regular consulting rate. I also receive a share of the profits of Auctionomics Inc., which has, in this matter, provided research support and assisted in the preparation

of this report under my direction and supervision. My compensation is not contingent on my findings, the testimony I may give, or the outcome of this case.

II. SUMMARY OF OPINIONS

A. Google's Programs Benefited Its Customers

14. Plaintiffs allege that "Google's exclusionary conduct has caused a wide range of anticompetitive effects, including higher prices, reduced output, lower quality services, reduced innovation, the exit of rival firms, and foreclosed entry in the relevant antitrust markets[.]"[10] After analyzing the Google practices listed in Paragraph 12, I find instead that these practices represent competition on the merits, providing benefits to Google's customers: its advertisers, publishers, or both.

15. I now provide brief summaries of my analyses of Google's programs and their effects, with details and supporting evidence to appear in later sections of this report.

   a. Google Ads' bid optimization programs—buy-side DRS, Bernanke, Global Bernanke, and Alchemist—optimized bids into the AdX auction to increase the total value of impressions won by Google Ads advertisers. Whenever these programs increased Google Ads' win rate (increasing the number of impressions won by Google Ads advertisers), they also increased the surplus enjoyed by Google Ads advertisers. Google Ads' experiments suggest that its bid optimization programs also benefited publishers in the form of increased revenues and a reduction in the number of unsold impressions, expanding output. Bernanke achieved those goals without turning AdX into a "third-price auction," as alleged by Plaintiffs.[11] While the specific details of Project Bernanke were not

[10] Fourth Amended Complaint ¶ 502.

[11] Fourth Amended Complaint ¶ 299 ("As addressed below, Google's secret Bernanke program surreptitiously switched Google's AdX exchange from a second-price auction to a third-price auction on billions of impressions per month."). *See also* Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 359 ("In Project Bernanke, participants believed

communicated to advertisers or publishers, the strategies of bidders in auctions are routinely kept confidential for the benefit of those bidders. Buy-side tools of other ad tech companies used similar bid optimization programs, and I am not aware of cases in which the bidding algorithms of those tools are made public.

b. Project Bell benefited Google Ads advertisers by protecting them from a publisher tactic called multi-calling, which would otherwise reduce their advertiser surplus. Multi-calling involves calling AdX multiple times for the same impression. Project Bell also benefited non-multi-calling publishers who might otherwise have received lower bids from advertisers seeking to protect themselves against multi-calling. Project Bell did not "punish" publishers who partnered with Google's competitors or turn AdX into a "third-price auction," as alleged by Plaintiffs.[12] Buy-side tools of other ad tech companies also modified bids in response to multi-calling.

c. Projects Poirot and Marple benefited advertisers using DV360 and Google Ads, respectively, by optimizing their bids to prevent advertisers from overpaying when auctions were not second-price auctions. Optimal bidding strategies can sometimes be complex to compute, so by automating this bid optimization, Poirot and Marple made bidding simpler and more profitable for Google's advertiser

they were in a second-price auction, but it was essentially a third-price auction, with the publisher receiving the third-highest bid, the advertiser paying the second-highest bid, and Google pooling the difference to manipulate other auctions.").

[12] Fourth Amended Complaint ¶¶ 311 ("If a publisher does not give preferential access to AdX, then Bell would drop their auctions from second- to third-price auctions[.]"), 557 ("Bell punished publishers who did this by dropping second price bids returned to publishers that had not enabled Dynamic Allocation (or otherwise ranked it near the top of their waterfall)").

customers. Buy-side tools of other ad tech companies used similar automatic bid shading programs for non-second-price auctions.

d. Project Elmo benefited Google Ads and DV360 advertisers by ensuring that their budgets were not depleted too quickly as a result of multi-calling by publishers and bid duplication by exchanges. Bid duplication occurs when an exchange sends multiple bid requests for the same impression in an attempt to elicit higher bids from one of those calls. By blocking the harmful effects of bid duplication, Elmo reduced spending on exchanges engaged in that practice and increased spending on other exchanges, while benefiting advertisers by spending their budgets more effectively. Project Elmo did not treat exchanges participating in header bidding differently from non-header bidding exchanges, as alleged by Plaintiffs.[13]

e. Dynamic Allocation (DA) benefited publishers by introducing real-time auctions that allowed them to sell impressions on AdX when the AdX offer was larger than the publisher's expected price from any other demand source. DA also benefited advertisers, allowing them to purchase impressions on AdX while paying *only* the amount they needed to bid to win impressions, and not more. Supply-side tools of other ad tech companies developed ad allocation mechanisms similar to DA.

f. Enhanced Dynamic Allocation (EDA) benefited publishers by allocating impressions between direct deals and remnant demand in a way that increased publisher revenues without compromising publishers' ability to fulfill direct

[13] Fourth Amended Complaint ¶ 403 ("[Elmo] decreased overall ad spend on any exchange that it suspected to meaningfully engage in header bidding.").

contracts with advertisers. EDA also expanded output by reducing the number of unsold impressions and helped advertisers to win the impressions they valued most. Supply-side tools of other ad tech companies implemented similar programs to optimize between direct deals and remnant demand.

g. The so-called "last look" of AdX over header bidding was not a Google-designed program but a side effect of the way that some publishers configured header bidding to integrate with Google's ad server. Publishers who used header bidding could benefit from offering AdX the chance to bid on inventory because it allowed them to earn higher revenues and integrate the other services that Google Ad Manager provided. The so-called "last look" did not create an inherent advantage for AdX bidders. Non-Google ad servers also used ad allocation mechanisms resulting in a similar "last look."

h. Reserve Price Optimization (RPO) increased publisher revenues and simplified revenue optimization for publishers in the AdX second-price auction. RPO did not alter the auction format, and it never unsealed bids received in an auction to set the floor price for that auction. Because publishers could adjust floor prices on the basis of historical data before and after the introduction of RPO, a surplus-maximizing bidder would need to account for the possibility that future floor prices would change in response to their bids both when RPO was in place and before.

i. Sell-side Dynamic Revenue Sharing (DRS) benefited publishers by allowing them to sell more inventory and increase their total revenues from the sale of

impressions. Sell-side DRS also expanded output by reducing the number of unsold impressions, allowing advertisers to win more inventory. Google's intention to use auction optimizations like sell-side DRS was transparently communicated to auction participants on Google's Help Center pages. Non-Google exchanges and supply-side intermediaries implemented similar programs to increase the win rates of their advertisers.

j. Open Bidding benefited publishers by allowing them to incorporate bids from non-Google exchanges into an auction among bids from multiple auctions—an "auction of auctions"—leading to higher auction revenues without the drawbacks of alternative approaches, including header bidding. Open Bidding also benefited advertisers on competing exchanges by increasing the total inventory open to competition from non-AdX bidders. The transition to the Unified First Price Auction allowed the auction of auctions to replicate the result of a first price auction in which all bidders are treated equally.

k. Unified Pricing Rules (UPR) allowed publishers to set floor prices natively in Google's supply-side platform (Google Ad Manager), but required those floors to apply uniformly across exchanges and demand sources. UPR protected multi-homing advertisers from price-fishing, a tactic in which some publishers would call bidders on different exchanges using different floor prices to induce them to make an unnecessarily high bid in at least one exchange. Without UPR, publishers would have incentives to engage in price-fishing, making coordinating bids from different channels more difficult for advertisers. Advertisers' likely responses—using fewer channels or reducing all their bids—could reduce

efficiency and harm publishers who did not engage in price-fishing. Non-Google ad tech intermediaries also introduced rules requiring floor prices to be uniform across bidders.

B. Analyses of Google's Programs by Plaintiffs and Their Experts Are Marred by Omissions and Errors

16. My opinions about Google's auction practices differ from those of Plaintiffs and their experts for four key reasons. *First*, Plaintiffs and their experts systematically overlook or understate the significant benefits that each of Google's auction practices confer on its advertiser and publisher customers. *Second*, Plaintiffs' experts exaggerate the challenge for multi-sided platforms of balancing the interests of buyers and sellers and omit the benefits that Google's integrated structure provides for its advertiser and publisher customers. *Third*, the analyses of Google's auction conducted by Plaintiffs' experts underestimate or understate the incentives for publishers and advertisers to optimize to improve their returns from display advertising. *Fourth*, those analyses underestimate or understate the prevalence and effectiveness of experimentation for optimizing returns. These omissions and errors lead Plaintiffs and their experts to false conclusions that Google's auction programs were anticompetitive and/or deceptive. Correcting these errors, a different explanation emerges as more consistent with the design of the technical programs at issue: each program helped Google compete on the merits for customers by providing greater benefits to advertisers, publishers, or both groups, including expanding output by enabling more matches to be made between advertisers and end users viewing publisher inventory.

17. I now discuss these four categories of omissions and errors individually, with additional details and supporting evidence to appear in later sections of this report.

1. Plaintiffs' Experts Overlook or Understate the Benefits of Google's Programs for Its Customers

18. Plaintiffs' experts focus their analyses on the effects of Google's programs on competitors, but in doing so, they largely overlook the significant benefits of Google's programs for its advertiser and publisher customers. Each of the Google programs I have studied benefits advertisers, publishers, or both, and, taken together, they have been output-expanding, that is, they have fostered the significant growth of online display advertising. Designing products that benefit consumers is the essence of competition on the merits, and the demand for Google's well-designed products contributed to the industry's (and the company's) growth.

19. To assess the effects of Google's auction programs on its advertiser and publisher customers, one must consider those programs in the historical context of the time of their introduction. To assist in that historical assessment, I provide a timeline of the introduction of Google's auction programs in Figure 1. The time period covered by Figure 1 was one marked by rapid technological development and growth of online display advertising. When each new Google program was introduced, it raised advertiser surplus or publisher revenues or both, and it sometimes created a standard on which future improvements could be built.

20. In several cases, Plaintiffs' and their experts' assessments of Google's conduct overlook this historical context. I provide a non-exhaustive list of examples from Plaintiffs' experts below:

   a. Plaintiffs' experts allege that publishers were harmed because Dynamic Allocation (DA) integrated real-time bids only from AdX,[14] but these allegations ignore the fact that real-time bidding was introduced into DA during the nascency of ad exchanges as a technology. At that time and in that context, industry participants viewed the main challenge not as exchange interoperability but "driving adoption," because "[t]he exchange represents a rather significant shift in how we typically transact, so adjusting to that for both buyer [and] seller takes some time."[15] DA eased this transition by making DA's design compatible with publishers' previous ad sales configurations, and Google drove adoption of the ad exchange technology partially through its integration of demand from Google Ads (then AdWords).[16] By overlooking this historical context, Plaintiffs ignore both the challenges that Google overcame in developing DA and the major benefits DA created for publishers.

[14] *See, e.g.*, Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 19 ("Google's Dynamic Allocation and Enhanced Dynamic Allocation distorted the playing field in its favor because Dynamic Allocation was solely granted to AdX and not competing exchanges."); Expert Report of J. Gans (Jun. 7, 2024), at ¶ 569 ("DA allowed AdX, and only AdX, to compete in real-time against all non-guaranteed inventory [...] Google understood that publishers were harmed by this feature of DA and that Header Bidding was the result of publishers seeking better prices.").

[15] Email from [redacted] to adx-updates@google.com, "FW: comments from industry players on AdX 2.0 on AdExchanger.com this evening" (Sep. 22, 2009), GOOG-AT-MDL-B-003180112, at -114.

[16] *See, e.g.*, Email from [redacted] to adx-updates@google.com, "FW: comments from industry players on AdX 2.0 on AdExchanger.com this evening" (Sep. 22, 2009), GOOG-AT-MDL-B-003180112, at -114 ("This will only work if Google can bring significant demand for display inventory into the system. They can maximize the demand by accepting bids from AdWords and also via networks through the AdX 2.0 channel.").

b. Plaintiffs' experts allege that Enhanced Dynamic Allocation (EDA) reduced the revenue that publishers could expect to earn from direct deals.[17] But Plaintiffs' experts overlook the fact that, during a period in which real-time bidding was rapidly growing in popularity, EDA helped publishers solve the difficult problem of how to allocate display inventory between direct deals and remnant demand channels. EDA solved this important problem by making it easier for publishers to monetize their inventory using *both* direct deals and remnant demand, and, as I discuss in Section IX, EDA led to revenue increases for publishers, without reducing the performance of direct deal advertising.

c. Plaintiffs' experts allege that Project Bernanke harmed non-Google exchanges and ad buying tools and created inefficiencies "by enabling lower-value advertisers to win impressions instead of higher-value ones."[18] Logically, the truth of this allegation depends not just on Google's strategy, but also on the bidding strategies of non-Google buying tools. What Plaintiffs' experts overlook is that, when Bernanke was introduced, many buying tools were submitting just a single bid into the Google AdX second-price auction. Such a one-bid policy may incentivize an advertiser or buy-side tool to submit bids into the AdX auction

[17] *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 137 ("In my opinion, Enhanced Dynamic Allocation likely led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers, which would in turn decrease the revenue earned by publishers via direct deals.").

[18] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 181; *see* Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.D.1 ("Under Projects Bernanke and Global Bernanke, GDN increased its revenue at the expense of non-Google ad-buying tools[.]"), Section VIII.E.1 ("Projects Bernanke and Global Bernanke did not benefit GDN advertisers, but decreased win rates for advertisers using non-Google ad buying tools[.]"); Expert Report of J. Gans (Jun. 7, 2024), at Section VIII.B.2 ("Bernanke harmed competition in the market for ad buying tools for small advertisers[.]"). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.C.2 ("Projects Bernanke and Global Bernanke can lead to a reduction in ad quality[.]"). Professor Weinberg's arguments only follow "if GDN advertisers tend to display lower quality ads," *id.* at ¶ 247, but he offers no justification for why this presupposition is true. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 759 ("Bernanke enables lower-quality ads to be transacted and displayed on publishers' properties.").

higher than the advertiser's true value for the impression.[19] Given that possibility, if Google Ads had continued to bid its advertisers' values into the AdX second-price auction, those advertisers would have been *disadvantaged* relative to non-Google Ads advertisers. Without a bid optimization strategy like Bernanke for Google Ads, the allocation of advertising inventory would have been inefficient for the *reverse* of the reason identified by Plaintiff's experts: non-Google Ads advertisers with lower values could win impressions at the expense of Google Ads advertisers with higher values.

d. Plaintiffs' experts claim that publishers had "historically" set different floor prices for different demand sources in order to preference some exchanges and that UPR "removed a key tool used by publishers to maximize the yield on their inventory,"[20] but these claims are misleading because they overlook the historical context. As I explain in Section VIII and Section XIV in this report, for publishers to maximize revenue under a sequential allocation process like the waterfall and DA, they needed to set floor prices that depend on each demand source's order in the waterfall, with later demand sources generally having lower floors. Because UPR was introduced together with the Unified First-Price Auction (UFPA), when

[19] For example, consider a buying tool that submitted a single bid into the AdX second-price auction on behalf of a group of advertisers, with that bid equal to the highest value reported by the advertisers in that group. If the buying tool charged its advertiser customers the clearing price of the auction, then that policy would lower each advertiser's expected payment for any bid submitted (compared to a counterfactual policy of submitting two bids equal to the values of their two highest-value advertisers) because the second bid submitted by the buying tool can only increase the clearing price of the auction and thus the price charged to the advertiser. Thus, the result of the one-bid policy is a reduction in the expected price paid for each possible bid, which creates an incentive for an advertiser to report higher bids to the buying tool (higher even than its value for the impression) to try to win additional impressions at lower average prices.

[20] Expert Report of P. Pathak (Jun. 7, 2024), at ¶¶ 157 ("Historically, publishers set higher reserve price floors for AdX to account for the perceived lower ad-quality of impressions served through AdX and increase diversity of demand sources."), 159 ("By eliminating publishers' ability to set differential price floors across exchanges, Google removed a key tool used by publishers to maximize the yield on their inventory and ensure acceptable quality advertisements were displayed on their web pages.").

the waterfall was replaced by a single auction, this motive for setting unequal floor prices was eliminated. The rest of this allegation, which suggests that UPR made it difficult for publishers to preference non-Google exchanges, is again misleading because publishers have better means of preferencing exchanges, including post-auction discounts, as discussed in Section XIV. UPR also served an important positive function: it reduced the harm to advertisers of price-fishing strategies that publishers might find to be individually profitable but that would raise costs for advertisers and undermine their trust in the marketplace.

HIGHLY CONFIDENTIAL – SUBJECT TO PROTECTIVE ORDER

Figure 1: Timeline of Google's Product Evolution

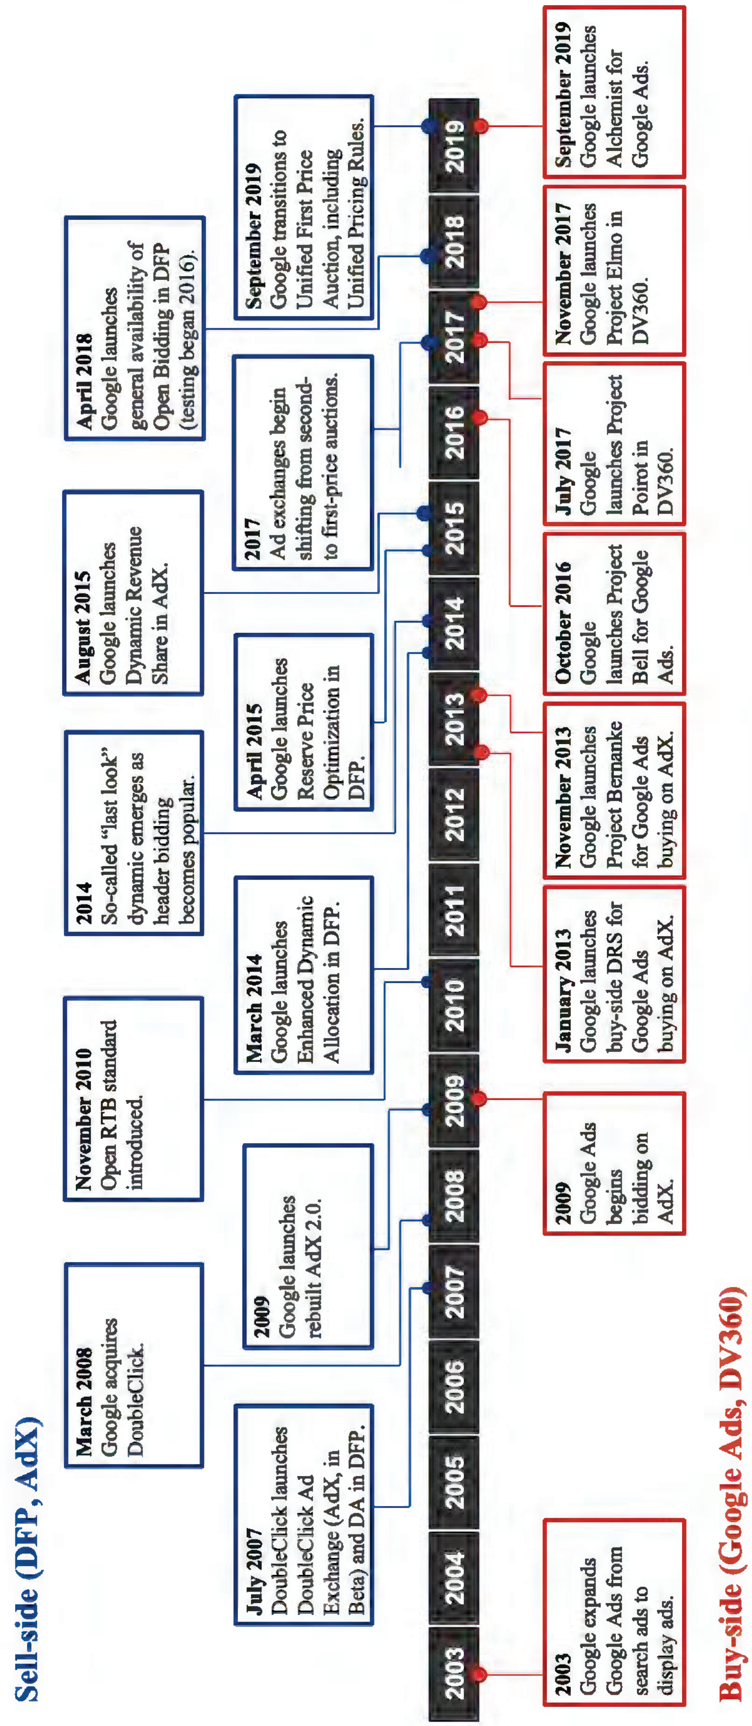

2. Plaintiffs' Experts Ignore the Benefits from Google's Business Model that Balances the Interests of Advertisers and Publishers

21. Professor Pathak asserts that "[b]uyers and sellers in marketplaces have opposing interests"[21] and that "[b]ecause Google is involved with all […] entities, it has an inherent conflict of interest."[22] But this characterization of marketplaces in general and Google's business model in particular is incomplete and omits the economic benefits arising from integration.

22. A matching market does not create a zero-sum game: buyers and sellers in matching marketplaces have a mix of aligned and conflicting interests. When a buyer's value for an impression exceeds the value of a seller's cost, *both* parties can benefit from the sale of the impression. Both sides can benefit when a well-designed marketplace improves match quality or facilitates more mutually beneficial transactions. Google, whose platform is paid based on trades between buyers and sellers, has an interest in growing the number and value of output-enhancing transactions on its platform.

23. Unlike an intermediary representing just one side of an industry, platforms like Google are incentivized to account for *externalities* that occur among participants on the platform. For example, where an intermediary representing sellers alone might be incentivized to engage in multi-calling (in which a publisher calls the same bidders

[21] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 17 ("Buyers and sellers in marketplaces have opposing interests. A buyer wishes to pay less for an impression, while a seller wants to receive more. The marketplace operator wishes to maximize trading volume and steer traffic to its exchange over competing alternatives.")

[22] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 17 ("Because Google is involved with all three entities, it has an inherent conflict of interest. Maximizing the interests of one type of participant may harm the interests of another type of participant. Google's conduct in the Ad Tech Stack results from conflicts of interest due to being involved in all sides of digital advertising transactions."). *See also* Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 244 ("This consolidation of power presents potential conflicts of interest, particularly when a single entity oversees both the buy-side and sell-side platforms, as well as the exchange where transactions occur.").

multiple times to bid for an impression, possibly with different floor prices, harming both advertisers and other publishers and undermining the auction design), a platform representing *both* publishers and advertisers would account for the harms of such behavior on advertisers and would disincentivize that behavior. Professor Pathak contends that various elements of Google's challenged conduct would not have arisen but for Google's integrated business model.[23] But those claims overlook the fact that, for all of the programs that Plaintiffs and their experts claim to be anticompetitive, there are less integrated display advertising intermediaries with the same or similar features: see Table 1 below. This means that the challenged conduct cannot be explained solely as a result of Google's integration, as Professor Pathak contends.

24. While Professor Pathak's observation of conflicting incentives between buyers and sellers is exaggerated, it is true that buyers and sellers have conflicting interests regarding *prices*, which determine how any gains from each transaction are shared. But Google's business model balances those conflicting interests in a disciplined way that jointly benefits all participants on the platform. From its *publishers*, Google collects a contractually-determined average share of revenue that results from its auctions, and for its Google Ads advertisers, it charges *threshold prices*, that is, an advertiser pays the minimum bid required for it to win an impression given the Google Ads bidding policy. Together, these two policies determine the payments made to publishers and by

[23] *See, e.g.*, Expert Report of P. Pathak (Jun. 7, 2024), at ¶¶ 96 ("Google acts on its conflicts of interest by taking actions that are contrary to the principles of market design I outlined above which give rise to well-functioning marketplaces."), 135 ("Dynamic Allocation was motivated by Google's conflict of interest to use its ad server in service of its exchange."), 141 ("Google's initiatives to undermine Header Bidding reduced marketplace efficiency by skewing the allocations of impressions to favor AdX through Exchange Bidding and reducing the ability of other exchanges to compete in real-time. Absent the conflicts of interest arising from Google's suite of display advertising products, DFP would have no incentive to undermine a technology that would maximize value for its publisher customers."), 158 ("UPR is also motived by Google's conflict of interest to use the DFP ad server to give preferential access to the AdX exchange.").

advertisers. Taking those pricing policies as given, Google offers tools to help advertisers and publishers optimize their returns from display advertising. The only remaining decisions that affect the platform's prices are ones that determine how Google collects its revenue share from advertisers and publishers. Those decisions also affect the *matching* between advertisers and publishers and thus the total value created by the platform. The net result is that Google is incentivized to optimize the total value of the matches on its platform, including via innovations that improve matching, such as DA, EDA and Open Bidding.

Table 1: Google's Competitors Developed Similar Product Features as Challenged Google Products

Feature	Challenged Google Product	Competitor(s) with a Similar Product Feature	Report Section
Buy-side: Bid Optimization Programs Varying Revenue Shares	Buy-side DRS, Bernanke, Alchemist	[redacted]	IV
Buy-side: Adjusting Bids for Multi-calling	Bell, Elmo	The Trade Desk (TTD); MediaMath (before 2023); [redacted]	V, VI
Buy-side: Adjusting Bids for Bid Duplication	Elmo	[redacted]	VI
Buy-side: Adjusting Bids to Auction Format	Poirot, Marple	[redacted]; The Trade Desk; MediaMath (before 2023); [redacted]	VII
Sell-side: Auctions with Floors Determined by Remnant Line Items	Dynamic Allocation	OpenX; [redacted]	VIII
Sell-side: Direct Competition of Guaranteed Contracts and Remnant	Enhanced Dynamic Allocation	[redacted] Magnite; Comcast's FreeWheel; OpenX; [redacted]	IX
Sell-side: Line Items Associated with Header Bidding Determining Auction Floors	So-called "Last Look"	Open X; [redacted]	X
Sell-side: Reserve Price Optimization	RPO	[redacted] Magnite (formerly Rubicon); [redacted]	XI
Sell-side: Varying Revenue Shares by Impression	Sell-Side DRS	[redacted]	XII
Sell-side: Requiring Uniform Reserve Prices by Demand Source	UPR	[redacted] Meta (code of conduct for partners); Xandr (best practices)	XIV

3. Plaintiffs' Experts' Analyses Underestimate the Importance of Incentives

25. Plaintiffs and their experts routinely underestimate or understate the ability of and incentives for advertisers and publishers to optimize their behavior when Google introduces or modifies its auction programs. It is my opinion that accounting for advertiser and publisher incentives to respond to auction programs is necessary to evaluate correctly the economic effects of these programs, and Plaintiffs' experts' analyses that fail to do so are unreliable.

26. In the subfield of economics known as Market Design, it is routine and necessary for academic papers that analyze the effects of marketplace rules, programs, and practices to pay careful attention to how incentives affect the behavior of marketplace participants.[24] As Professor Weinberg acknowledges, this canonical method of analysis is "commonly accepted by researchers and practitioners for the analysis of the market at hand, online display ads."[25] The same method is described and applied in academic work by three of the Plaintiffs' experts—Professors Weinberg, Pathak, and Gans[26]—but it has not been

[24] In the syllabus of my Stanford class on this subject, I explain that "'Market design' is the subfield of economics that studies how best to organize *decentralized* resource allocation systems taking especially careful account of how the rules of the system affect individual *incentives* and *choices*. 'Decentralized' means that the system relies sensitively on information sourced from individual participants." *See* Paul Milgrom, "Market Design" Syllabus, Stanford University (accessed Jan. 10, 2024), https://canvas.stanford.edu/courses/171062. Similarly, the course description of Professor Susan Athey's Stanford University course entitled "Topics in Market Design" announces that it "studies the design of organized markets, focusing on efficient organization and *the incentives created by market rules*. Applications include online auction markets […] ." *See* Susan Athey, "Economics 980: Topics in Market Design," Stanford University (accessed Sep. 27, 2023) (emphasis added), https://gsb-faculty.stanford.edu/susan-athey/economics-980-topics-market-design/. Professor Weinberg explains that he is an expert in the closely-related field of "Algorithmic Mechanism Design, which is the study of algorithms (such as ad auctions) that involve economic incentives (such as those of publishers, exchanges, ad buying tools, and advertisers)." *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 4.

[25] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 13.

[26] *See, e.g.*, Cai, Y., Daskalakis, C., & Weinberg, S. M. (2013). Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (pp. 618-627) ("*Mechanism design* is the problem of optimizing an objective subject to 'rational inputs.' The difference to *algorithm design* is that the inputs to the objective are not known, but are owned by rational agents who need to be provided incentives in order to share enough information about their inputs such that the desired objective can be

applied consistently across the analyses in their expert reports. For theoretical predictions, a common approach to studying incentives in economic analyses of auctions is to investigate *equilibrium* whenever that is possible. "Equilibrium" means that all agents choose actions to maximize their benefits of participating in the auction and that they make accurate forecasts about other agents' choices.[27]

27. Although accounting fully for incentives is the standard and indispensable benchmark for evaluating the effects of Google's auction-related programs on its customers, it can be useful to supplement that analysis with one that focuses on what happens over shorter periods of time, when self-interested participants have not yet fully identified the new incentives and learned how best to adjust their actions. To study those shorter-lasting effects, I also incorporate analyses that assume behavior continues unchanged for a period after Google introduces a new program.

28. It is my opinion that Plaintiffs' experts' analyses that omit participants' incentives are unreliable, falling short of ordinary professional standards. When Professor Weinberg characterizes floor price optimizations by publishers as applying only to publishers who are "sophisticated" or "clever," he offers neither evidence nor logic to justify his

optimized."); Kojima, F., & Pathak, P.A. (2009). Incentives and stability in large two-sided matching markets. *American Economic Review, 99*(3), 608-627 ("Under some regularity conditions, we show that the fraction of participants with incentives to misrepresent their preferences when others are truthful approaches zero as the market becomes large."); Gans, J.S., & Holden, R.T. (2022). Mechanism design approaches to blockchain consensus, National Bureau of Economic Research Working Paper No. w30189 ("The question we address in this paper is whether there are more efficient and more reliable ways to achieve truth in consensus by designing and encoding mechanisms. Mechanism design is the branch of economics that deals with creating incentives for self-interested agents with information not known to the designer to reveal that information truthfully and still be willing to participate in the relevant economic activity.").

[27] Bayes-Nash equilibrium is a standard solution concept of game theory. *See* Harsanyi, J.C. (1967). Games with incomplete information played by "Bayesian" players, I-III. Part I. The basic model. *Management Science*, *14*(3), 159-182; Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT Press, at 3, 11-14.

decidedly non-standard approach.[28] In reality, floor price adjustments require no more sophistication than is routine in economic decision-making. For example, just as someone selling a used car who has been offered $10,000 for the car by a dealership benefits by offering it for a private sale at a higher price (say $12,000), rather than at the same $10,000 price, a publisher who has been offered $3 for an impression by a header bidder benefits by offering it in its auction at a higher floor price, rather than at the same $3 price. No special sophistication is needed to understand that.

29. Similarly, advertisers and the intermediaries that represent them need no unusual sophistication to identify and respond to the incentives discussed in this report. For example, a bidder (on DV360, say) that bids higher than the floor price and observes that it is routinely being charged its bid by a non-Google exchange purporting to run a second-price auction needs only to look at that data to learn that the exchange was actually running a first-price auction and that it should shade its bid to avoid overpaying for impressions. Yet that bidder response is ignored by Professor Gans in the analysis supporting his conclusion that Project Poirot "reallocat[ed] revenue from rival exchanges to Google's own exchange"[29] Empirical evidence from online display advertising

[28] *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 159 ("If a sophisticated publisher instead cleverly sets Value CPMs as a function of header bids, then AdX might still infer information about the maximum header bid, and in particular certainly knows that the maximum header bid lies below its reserve."), ft. 165 ("The impact of Dynamic Allocation with sophisticated publishers who cleverly set Value CPMs is less clear-cut. On one hand, if sophisticated publishers only slightly inflate the Value CPM of the winning header bid, then the above conclusions continue to hold for exactly the same reasons. On the other hand, if sophisticated publishers significantly inflate the Value CPM of the winning header bid due to Dynamic Allocation and would not have set such an inflated reserve on AdX in absence of Dynamic Allocation, then the cost of this inflated reserve might outweigh the benefits highlighted above."), ft. 225 ("If a sophisticated publisher mildly inflates AdX's reserve specifically because of the Last Look advantage, or if advertisers bid similarly in these two cases, the same conclusions still qualitatively hold. If a sophisticated publisher drastically inflates AdX's reserve specifically because of the Last Look advantage, or advertisers drastically change their bids specifically due to AdX's Last Look advantage, the impact is less clear-cut and would require a complicated analysis weighing the benefits of Last Look versus the impact of an increased reserve and distinct bids.").

[29] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865.

auctions suggests that bidders *do* learn to respond to auction design changes over time, and eventually come to adopt nearly profit-maximizing strategies.[30]

30. That finding is no accident. Many advertisers contract with specialized intermediaries (such as advertising agencies) to perform similar optimizations for them,[31] and some advertisers and publishers even employ teams of engineers, economists, and marketing experts devoted to maximizing returns by finding all possible improvements in advertising yields. And even though advertisers and publishers do not in all cases perfectly optimize their returns from display advertising, it is common for them to conduct experiments and learn to adjust their bidding and/or floor pricing strategies, which tends to bring them closer over time to the equilibrium benchmark.

31. For many publishers and advertisers, there are vast sums of money at stake, and online display advertising makes up a significant fraction of their marketing revenue or spend. Moreover, the largest advertisers and publishers make up a significant fraction of the online display advertising industry, suggesting that the strategic sophistication of these agents is especially relevant to auction outcomes.[32] Academic research has often found

[30] *See, e.g.*, Goke, S., Weintraub, G. Y., Mastromonaco, R., & Seljan, S. (2022). Bidders' responses to auction format change in internet display advertising auctions. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4021221.

[31] *See, e.g.*, [redacted]

[32] For example, in 2022, more than [redacted]% of Google Ads' US web spending across AdX, AdSense, and non-Google exchanges was by advertisers spending more than [redacted] on the platform. For the same criteria, more than [redacted]% of DV360 spending was by advertisers spending more than $[redacted]. *See* GOOG-AT-MDL-DATA-000486626 to

that bidders with large stakes in auctions typically behave in line with their incentives, as predicted by equilibrium analysis.[33]

4. Plaintiffs' Experts' Analyses Underestimate the Role of Experimentation for Optimizing Returns

32. Plaintiffs and their experts claim that the optimization process for publishers and advertisers was compromised because "Google did not disclose their secret programs."[34] Plaintiffs' experts suggest that this alleged failure to disclose would mislead a "typical" advertiser or publisher, keeping it from optimizing its choices effectively.[35] In reality, processes for setting reserve prices and determining bids in auctions are routinely kept secret to prevent other auction participants from gaming these strategic choices. And

-8277 (Google Ads); GOOG-AT-MDL-DATA-000561263 to -420 (XBridge DV360). This result was generated using code/web_spend.py in my supporting materials, and the output is saved in code/logs/web_spend.txt.

[33] *See, e.g.*, Doraszelski, U., Lewis, G., & Pakes, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, *108*(3), 565-615; Hortaçsu, A., & Puller, S. L. (2008). Understanding strategic bidding in multi-unit auctions: A case study of the Texas electricity spot market. *The RAND Journal of Economics*, *39*(1), 86-114; Hortaçsu, A., Luco, F., Puller, S. L., & Zhu, D. (2019). Does strategic ability affect efficiency? Evidence from electricity markets. *American Economic Review*, *109*(12), 4302-4342.

[34] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 197 ("Finally, Google decreased transparency by deceptively changing auction rules and bidding rules, as I discuss in Section XII. Google did not disclose their secret programs. Had publishers and advertisers known about these programs, they would have the opportunity to adjust their behavior."). *See also* Fourth Amended Complaint ¶ 328 ("Google continued to mislead publishers and advertisers about the program and withheld critical information that the parties could have used to make an informed decision about the program."); Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 250 ("Moreover, all publishers likely would have changed their behavior if they knew about Projects Bernanke and Global Bernanke by raising their reserve prices.").

[35] *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 114 ("On one end, a 'typical' publisher may set parameters according to their ad server's suggested text without developing a detailed understanding of how those parameters are used. At the other end, a 'sophisticated' publisher may fully digest all available documentation and aim to optimize parameters based on their use case, ignoring suggested text. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation. […] On one end, a 'typical' advertiser may trust their ad buying tool to optimize on their behalf and input correct information whenever requested (i.e., a 'typical' advertiser would simply input their correct value for an impression when asked). At the other end, a 'sophisticated' advertiser may fully digest all available documentation and aim to optimize inputs to their ad buying tool based on how these inputs are used, ignoring the ad buying tool's recommendations. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation.").

even for those programs for which details are not disclosed at all, advertisers' and publishers' routine data analysis and experimentation with bids and floor prices are typically sufficient for them to identify optimal strategies. In practice, evidence suggests that ad tech intermediaries, in-house marketing teams, ad agencies, and publishers rely heavily on feedback and experimentation to optimize their performance:

a. *Advertisers* leverage key performance indicators to guide their campaign strategies on buy-side tools and bid effectively. Rather than calculating bids themselves, advertisers delegate many of the details of bid optimization to specialized buy-side tools[36] or agencies,[37] while optimizing their campaign parameters to achieve higher click-through rates, conversion rates, or return on ad spend. [redacted]

[36] For instance, [redacted]; Jenna Naylon, "Media Planning and Buying Strategies: The Importance of Test & Learn," MatrixPoint (accessed Jul. 17, 2024), https://www.thematrixpoint.com/resources/articles/media-planning-and-buying-strategies-the-importance-of-test-learn ("Applying a systematic approach to 'testing' (running controlled 'what ifs') varying aspects of media campaigns and 'learning' from the results to refine strategies, drive better results, and maximize marketing ROI. The core principle is to iterate and optimize based on data and insights gathered from continuous testing. Successful test and learn strategies include setting clear objectives, selecting appropriate key performance indicators (KPIs), allocating resources effectively, conducting A/B testing or multivariate testing, and implementing robust measurement and analysis procedures.").

[37] Both large and small advertisers consult with specialized agencies to achieve further sophistication. *See, e.g.*, [redacted]

[redacted][38] [redacted]"[39]

b. *Buy-side tools* depend on considerable experimentation, employing various learning algorithms to optimize bids on behalf of advertisers. Through experimentation, these algorithms can automatically and almost immediately adapt to changes in the environment. For example, an internal Google document found that certain buyers adjusted their bids in response to RPO, with the author noting, "[f]or RPO, some buyers are changing their bids […] some bid higher, some bid lower."[40] Similarly, documents suggest that Criteo detected experiments related to DRS through Criteo's own experiments that it ran for the purposes of bid optimization.[41]

[38] [redacted]

[39] [redacted]

[40] *See* "Display Ads Research Meeting Notes" (Jun. 19, 2017), GOOG-TEX-00831373, at -378. *See also* "AdX Managed Reserves" (Feb. 10, 2017), GOOG-DOJ-03643284, at -287 ("We also have evidence of third-party buyers bidding lower when we send RPO reserves.").

[41] Email from [redacted] "Re: [criteo] Dynamic Pricing on eBay UK" (Feb. 2016), GOOG-DOJ-15426012, at -018, ("[Criteo employee:] Our team has found data suggesting that there is dynamic pricing on eBay in the UK. It looks like dynamic pricing starts from PMP [private market place] floors and follows a (floor+bid)/2 line.").

c. *Publishers* frequently experiment to decide floor prices. A [redacted] There are supply-side intermediaries who assist in this process, as well, including Clickio, which offers floor price optimization tools and other services for publishers to "guarantee maximum revenues."[44]

33. Furthermore, Google's tools make such experimentation easier for its customers. For advertisers using Google Ads, an "Experiments" feature allows advertisers to conduct A/B tests on different campaign settings, including bidding strategies.[45] Advertisers using DV360 can use the "Insights" feature to conduct A/B tests on campaigns.[46] Google Ad Manager also contains functionality allowing publishers to experiment with various aspects of their properties, including floor prices, header bidding, and blocking.[47] Nor is

---

[42] "[redacted]

[43] [redacted]

[44] Clickio, "Monetization" (accessed Jul. 7, 2024), https://clickio.com/monetization/ ("Automated price floor optimization can boost earnings by 20-120%, while also reducing the number of ads on a page, leading to a better user experience.").

[45] *See* Google, "Test Campaigns with Ease with Ads Experiments," Google Ads (accessed Jun. 26, 2024), https://ads.google.com/intl/en_us/home/tools/experiments.

[46] Luke Hedrick, "Get actionable measurement with Display & Video 360's Insights module," Google Marketing Platform (Nov. 15, 2018), https://blog.google/products/marketingplatform/360/get-actionable-measurement-display-video-360s-insights-module/.

[47] Google, "Run a manual experiment," Google Ad Manager Help (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/9799933?hl=en#zippy=%2Cuser-messages%2Cnative-ad-style%2Cunified-pricing-rules%2Cyield-groups%2Cheader-bidding-trafficking-for-prebid.

Google alone in offering such tools: other advertising intermediaries also offer experimentation tools for their advertiser and publisher customers.[48]

34. By ignoring the evidence that publishers and advertisers experiment to optimize their returns from display advertising, Plaintiffs and their experts necessarily *underestimate* the amounts that Google's customers can earn from these programs. Across the wide range of programs I study in this report, I show that correcting this failure by Plaintiffs reveals significant benefits for advertisers and publishers from Google's auction features.

[48] *See, e.g.*, Microsoft, "Discover the possibilities with experiments," Microsoft Advertising (accessed Jun. 26, 2024), https://help.ads.microsoft.com/apex/index/3/en/56908; Magnite Team, "Magnite Unveils New Demand Manager Feature Powered by Machine Learning to Help Publishers Earn Incremental Revenue" (Oct. 5, 2023), Magnite, https://www.magnite.com/press/magnite-unveils-demand-manager-machine-learning-feature/; The Trade Desk, "Conversion Lift: What it is, how it works, and best practices" (Aug. 23, 2023), https://www.thetradedesk.com/us/resource-desk/best-practices-for-better-conversion-lift.

III. BACKGROUND AND ECONOMIC FRAMEWORK

A. The Key Economic Features of Online Display Advertising

35. Marketplaces are meeting places for trade. Buyers with wants and needs meet sellers of goods or services that can fulfill those wants and needs. When a buyer's value for a good or service exceeds a seller's cost, *both* parties can benefit from trade. Marketplaces provide social value when they facilitate additional opportunities for mutually profitable trade among economic agents.

36. In online display advertising, the key participants include website **publishers**, who sell ads on their websites (also known as **impressions**);[49] the **advertisers**, who buy impressions; and the **end users**, who may view and/or interact with the ads. Publishers use advertising sales to fund their production of internet content, which is often free or subsidized for end users. Advertisers can have many and diverse goals for their advertising campaigns, but they generally seek to increase the probability that an end user engages with their product, service, or message. End users experience benefits and costs from display ads: they may benefit directly from the information conveyed by the ad or indirectly from the ability to consume internet content at a reduced cost (*e.g.*, without facing a paywall), and they may experience costs in the form of the annoyance caused by unwanted advertisements or slower load times for web pages. In this report, because the

[49] In this report, I use the word "impression" to refer to an opportunity for a display advertisement that is offered for sale by a web publisher. This is a broader definition than the technical definition of impression typically used in the industry, which requires that the impression opportunity is successfully allocated to an advertiser and the associated advertisement loads successfully on the end user's browser. For example, Google Ads defines an impression as "[h]ow often your ad is shown[:] An impression is counted each time your ad is shown on a search result page or other site on the Google Network." Google, "Impressions: Definition," Google Ads Help (accessed Dec. 16, 2023), https://support.google.com/google-ads/answer/6320?hl=en.

costs and benefits to end users are difficult to measure and are not the main subject of Plaintiffs' allegations, I focus on the costs and benefits to publishers and advertisers.

37. Occasionally, I will refer to **economic welfare**, by which I will mean the total value to advertisers, minus any costs incurred by publishers and intermediaries. This definition of economic welfare is independent of the prices paid by advertisers for impressions and the fees charged by intermediaries, with these factors determining how economic welfare gets split among advertisers, publishers, and intermediaries.[50] A change in practices is said to **increase efficiency** if it increases economic welfare. For example, a change in matching procedures that leads to assigning impressions to advertisers with higher values increases efficiency and may benefit *both* publishers and advertisers.

38. One of the key challenges facing an advertiser is to identify the appropriate audience for its advertising campaign. The probability that any single ad successfully influences the end user (for example, to click on the ad, buy a product, sign up to an email list, or vote for a political candidate) is typically low.[51] At the same time, the potential benefits to the advertiser from a successful interaction, called a **conversion**, can be significant.

39. In economics parlance, online display advertising is a **matching market** because the value of an impression to an advertiser typically depends on various factors, including the ad shown, the identity of the end user, and the context of the ad. Advertisers are willing

[50] Economic welfare can equivalently be defined as the sum of the advertiser surplus earned by advertisers (see the definition of advertiser surplus in Paragraph 53 below), plus the profits earned by publishers, plus the profits earned by any intermediaries.

[51] For example, according to the Google Ad Manager Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000276098 to -001116097, about [redacted]% of US impressions won by bidders with transaction types "Open Bidding" or "Open Auction" resulted in a click. This result was generated using [redacted] in my supporting materials, and the output is saved in [redacted]

to pay more to show their ads to end users for whom that ad is more likely to be more relevant. For instance, a restaurant in Dallas, Texas is unlikely to derive much value from advertising its specials to an end user in San Francisco, but may have a high value for advertising them to end users in Dallas, especially those who dine out frequently or have visited that restaurant in the past. The same restaurant might derive more benefit from an advertisement next to a restaurant review in The Dallas Morning News online than from an ad on the sports page. Other advertisers, such as Ticketmaster, might have the opposite preference. Improved matching that places higher value ads for each impression can increase both the price paid to publishers and the profits earned by advertisers.

40. Another key economic feature of online display advertising is that ad impressions on webpages are quickly **perishable**: they must be allocated within fractions of a second of the end user's arrival to maximize the chance that the ad will be noticed. The perishability of impression opportunities distinguishes the matching problem in online display advertising from that faced on other online matching platforms (*e.g.*, real estate platforms like Zillow or dating websites like eHarmony) in which there is typically much more time for counterparties to consider and finalize the details of their matches.

41. These two characteristics—that good matching is key to creating value and that impressions are perishable—distinguish online display advertising platforms from commodity exchanges and from trading platforms for securities like stocks and bonds. In the sale of securities, buyers and sellers often care little about the identity of their trading partners. Partly as a result of that, trading platforms for securities are most often anonymous. In addition, securities are less perishable, with investors sometimes postponing their trading for days or longer if they are unsatisfied with current prices.

B. Intermediation in Online Display Advertising

1. How Intermediation Can Improve Matching

42. The challenges associated with matching impressions to end users encourage **intermediaries** to offer increasingly effective services. These services may increase economic welfare by allowing advertisers and publishers to quickly identify more and better matches and/or reducing the costs of participation (sometimes called **market frictions** or **transaction costs**). Such intermediaries typically share in the benefits of increased economic welfare by charging fees for their services.

43. Common approaches taken by intermediaries that increase economic welfare include:[52]

a. *Making it easier for publishers and advertisers to transact:* Intermediaries can make participation easier by performing technical and strategic computations that help publishers to manage their inventories or advertisers to evaluate and bid on millions of impressions. Without such intermediation, each buyer and seller might incur the costs of developing processes, performing computations, and making the inevitable errors that result from the complexity and frequent changes of an industry that is evolving and growing. In my own auction consulting in this industry and others, I have emphasized the importance of making bidding easier to encourage participation and promote value creation.[53]

[52] *See, e.g.*, Roth, A.E. (2015). *Who gets what—and why: The new economics of matchmaking and market design*. Houghton Mifflin Harcourt, at 8-11.

[53] For example, in my published advice as consultant to the US Federal Communications Commission, I wrote that "the auction process needs to be simple and easy enough to encourage and facilitate the participation of a wide array of broadcasters […] [and] make it very easy for broadcasters to make optimal bids." Auctionomics and Power Auctions, "Incentive Auction Rules Option and Discussion," FCC (Sep. 12, 2012), https://docs.fcc.gov/public/attachments/FCC-12-118A2.pdf, at 2.

b. *Making participation safer:* Intermediaries often design and enforce rules to protect participants from being taken advantage of by other participants on the platform. A lack of safety can force participants to spend resources monitoring and strategizing to protect against unscrupulous behavior. Such expenses can be wasteful and discourage participation.

c. *Making the platform thicker:* Matching platforms work better when they are *thicker*, which for online ad platforms means more advertisers, more impressions, and better information.[54] Thicker matching platforms enable advertisers to find more suitable impressions and to spend their budgets more effectively. They may also enable publishers to attract more bidders and higher prices for their impressions. In these ways, thicker matching platforms improve economic welfare. For online advertising, intermediaries can promote thickness and create more valuable matches by exposing each bidder to a larger number of relevant impressions or providing the information that bidders need to evaluate those impressions.

[54] I note that this definition of thickness does not explicitly include the number of publishers. In online display advertising, the value of an advertisement is primarily driven by the match between an advertiser and an end user, so the most relevant measure of thickness is of advertisers and the number of impressions, not the number of publishers. Even a single publisher can have many end users and many impressions (*e.g.*, Facebook), so a platform matching many advertisers to a single publisher's impressions could also be thick. Academic research studying online advertising often define thickness by the number of advertisers competing per impression. *See, e.g.*, Levin, J., & Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review: Papers & Proceedings*, *100*(2), 603-07, at 607; Ye, Z., Zhang, D.J., Zhang, H., Zhang, R., Chen, X., & Xu, Z. (2023). Cold start to improve market thickness on online advertising platforms: Data-driven algorithms and field experiments. *Management Science*, *69*(7), 3838-60, at 3839 ("Throughout this paper, we use the market thickness to represent the average number of ads competing for user impressions on an online advertising platform."); D'Annunzio, A., & Russo, A. (2024). Intermediaries in the online advertising market. *Marketing Science*, *43*(1), 33-53, at 34 ("Advertising markets differ in their thickness (*i.e.*, the number of advertisers belonging to that market)").

d. *Making processes more efficient:* Fixing the set of participants and the information they receive, intermediaries can increase economic welfare by adopting technologies that enable more offers to be directly compared, so that the allocation process can assign each impression to the participating advertiser who values it most highly.

e. *Reducing latency:* Intermediaries can reduce transaction costs by eliminating unnecessary and duplicative steps and increasing the speed at which transaction opportunities are identified and processed. This reduces the likelihood that an end user leaves the website before the ad is presented.

44. The perishability of impressions creates special challenges for intermediaries in online display advertising beyond those faced on other matching platforms. Matching platforms in other industries often help buyers and sellers identify potential matches but leave it to the parties themselves to finalize the details of their transactions (as on, for example, real estate platforms like Zillow, hotel booking platforms like Booking.com, or dating platforms like eHarmony). In contrast, online display advertising intermediaries typically need to automate the entire transaction to allow ads to be shown almost immediately. The advent of **programmatic advertising** was a response to this challenge, in which intermediaries create automated systems to analyze data, compare options, and make decisions to buy or sell impressions for publishers and advertisers, all in a few hundred milliseconds.

2. The Ad Tech "Stack"

45. To illustrate the various roles of online display advertising intermediaries, consider the following simplified depiction of one of the pathways in the programmatic advertising supply chain. First, an **end user** arrives at a publisher's website, which has an impression opportunity available to be filled by an ad. While the webpage is loading, a **call** can be made to the **publisher's ad server**, which is an intermediary that helps publishers manage their online display advertising inventory. The publisher's ad server may make an initial decision about the allocation of the impression, based on parameters chosen by the publisher. In some cases, the server might allocate the impression directly to an advertiser who has a pre-arranged contract with the publisher. In other cases, the publisher's ad server might make a call to one or more **ad exchanges** or **supply-side platforms (SSPs)**, which allow publishers to sell inventory using programmatic sales mechanisms, including auctions. These auctions typically involve calls to **demand-side platforms (DSPs)**, which are intermediaries that help advertisers purchase online display advertising inventory.[55] The publisher's ad server or the ad exchange might also make a call to one or more **ad networks***,* which buy online display advertising inventory directly from publishers or through other intermediaries on behalf of advertisers in their networks. After the winning advertiser is determined, the advertising content is transmitted to the end user's browser by the **advertiser's ad server***,* which is an intermediary that manages and stores advertisers' ads. This entire process is completed within a fraction of a second after the arrival of the end user at the publisher's webpage.

[55] I use the terms "buying tool," "buy-side tool," and "DSP" interchangeably in this report.

46. Collectively, these intermediaries are sometimes referred to as the **ad tech stack**, depicted graphically in Figure 2.

Figure 2: Simplified Depiction of the Ad Tech Stack

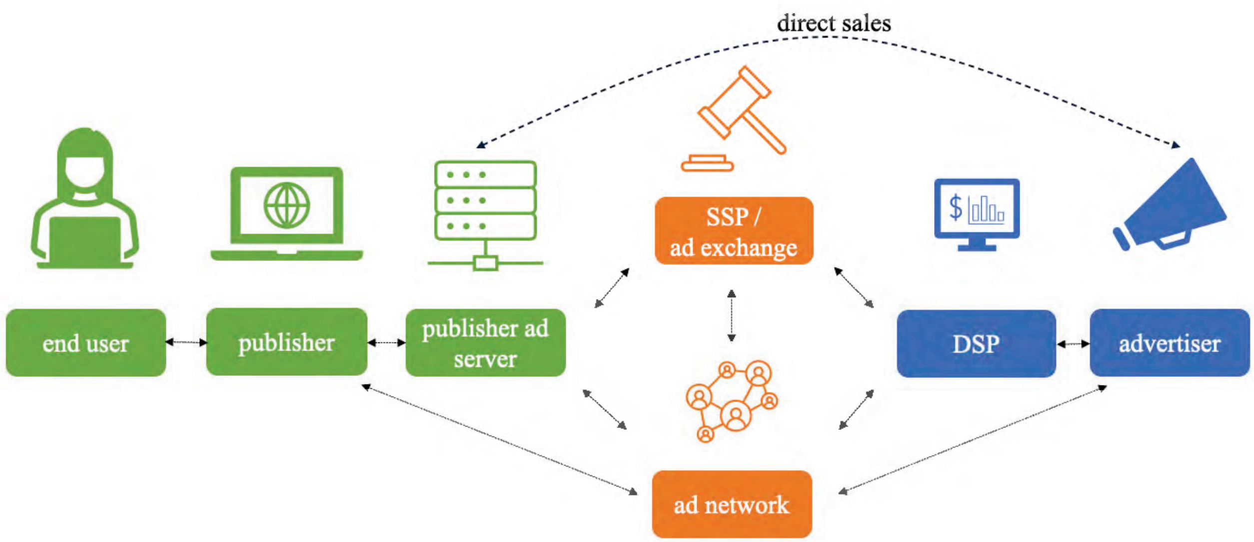

47. I emphasize that the above is a *simplified* description of the allocation process for *some* online display advertising impressions. Different publishers and supply-side platforms may use allocation processes different from the ones I described above, and advertisers may reach end users through different pathways. Advertising technology has evolved considerably over the past two decades and continues to evolve. Other important intermediaries, which may participate in the ad tech stack but are not included in my simplified description above, include **ad agencies**, which assist advertisers in planning and designing their advertising campaigns, and other intermediaries that assist advertisers and publishers, manage payments, and track user engagement. This description also omits **header bidding**, a technology that allows publishers to call demand sources directly, without first calling a publisher ad server. I discuss header bidding in detail in Sections X and XIII.

48. Another significant function of online display advertising intermediaries is the collection and management of data used by advertisers to identify end users and assess their values for display advertising inventory. One of the key tools used in this process is the **cookie**, which is a small text file stored in an end user's browser.[56] Cookies come in two main varieties: **first-party cookies**, allowing publishers to track user behavior on their own websites, and **third-party cookies**, allowing intermediaries to track user behavior *across* websites.[57] When a user first loads a website including content from an online display advertising intermediary, a third-party cookie is stored in the user's browser. After that time (or until that cookie is deleted in the end user's browser), the display advertising intermediary can track that user's activity across different websites that also include any content from that intermediary.[58] This allows the intermediary to develop a profile of an end user based on his/her browsing behavior, which can help advertisers assess the value of impressions associated with that cookie.[59] Because different display advertising intermediaries have different cookie information about end users, a process called **cookie**

[56] [redacted] (Redacted) (Sep. 29, 2023), GOOG-AT-MDL-C-000016753, at ¶ 14 ("In general, cookies are small data files stored on a web browser that can serve different functions.").

[57] Maciej Zawadziński & Mike Sweeney, "What is cookie syncing and how does it work?," Clearcode Blog (Jan. 31, 2024), https://clearcode.cc/blog/cookie-syncing/ ("Different Types of Cookies [...] First-party cookies are created by the websites we visit directly. [...] Third-party cookies, also referred to as tracking cookies, are collected not by the website, but by advertisers.").

[58] Maciej Zawadziński & Mike Sweeney, "What is cookie syncing and how does it work?," Clearcode Blog (Jan. 31, 2024), https://clearcode.cc/blog/cookie-syncing/ ("Third-party trackers can also track a user's behavior, such as the content they view on that website and the things they click on (e.g. products and ads). The trackers create third-party cookies and use them to display adverts to the user when they visit different websites. [...] Each time a user visits a website that contains ads (or third-party tracking tags), the browser sends an ad request to an advertising technology platform (e.g. a DSP).").

[59] *See, e.g.*, [redacted] (Redacted) (Sep. 29, 2023), GOOG-AT-MDL-C-000016753, at ¶¶ 16 ("[Cookie matching] allows an RTB participant, for example, to limit the bid requests they receive to those involving users that they previously interacted with, as determined by the presence of their cookies."), 39 ("Broadly speaking, these settings and controls will determine [...] whether cookies or other pseudonymous identifiers can be included (where they are available) which enables RTB participants to select ads based on information they may have on prior activity associated with the identifier.").

matching (or cookie syncing) is sometimes used to match (often imperfectly) cookies collected by one advertising intermediary to those collected by another, with this process aided by additional intermediaries, called **data management platforms**.[60] Google offers a cookie matching service to allow bidders in its display advertising auctions to match their cookies with Google's proprietary cookies, called Biscotti.[61] Advertisers often supplement data collected via third-party cookies with other first-party information about the end user to determine their value for displaying an ad.[62]

C. Information, Incentives, and Auctions

1. Auctions Aggregate Dispersed Information

49. One of the challenges associated with efficient matching and intermediation in online display advertising is that information about the value of an impression to different advertisers is typically dispersed and not directly observed by all advertisers, publishers, and intermediaries. An advertiser does not typically have all the information required to assess its value for showing an ad to the end user and needs to rely on information collected by publishers and intermediaries about the user's characteristics and browsing behavior. Meanwhile, a publisher seeking to sell online display advertising inventory

[60] Maciej Zawadziński & Mike Sweeney, "What is cookie syncing and how does it work?," Clearcode Blog (Jan. 31, 2024) https://clearcode.cc/blog/cookie-syncing/ ("An example of this would be mapping a user's ID from a demand-side platform (DSP) to a data management platform (DMP). This process is known as cookie syncing. [...] Cookie syncing works when two different advertising systems (aka platforms) map each other's unique IDs and subsequently share information that they have both gathered about the same user.").

[61] [redacted] (Redacted) [redacted], GOOG-AT-MDL-C-000016753, at ¶ 16 ("RTB participants may utilize their own cookies, in the same way Google uses a Biscotti. 'Cookie matching' allows RTB participants to match their cookies with Google's Biscotti for the same browser.").

[62] *See* [redacted] (Redacted) [redacted], GOOG-AT-MDL-C-000016753, at ¶ 39 ("Broadly speaking, these settings and controls will determine [...] whether cookies or other pseudonymous identifiers can be included (where they are available) which enables RTB participants to select ads based on information they may have on prior activity associated with the identifier.").

typically does not know the value of each impression to each potential buyer, or even the identities of all the potential buyers for a given impression. If advertisers and publishers had all of that information, auctions and many other services offered by intermediaries would be unnecessary because the publisher could simply sell each impression to the highest-value advertiser at a mutually agreeable price. Information processing is key to effective matching and pricing.

50. Auctions can perform a central role in this process: the auctioneer *distributes* certain information about the item being sold and receives information in the form of **bids** and uses those to allocate the item and determine payments. In this report, I focus on **sealed-bid auctions**, which are auctions in which bidders report their bids to the system just once at the start of each auction.[63] Because they can be resolved so quickly, sealed-bid auctions work well for online display advertising sales, where impressions are rapidly perishable.

51. In most auctions for a single item, the **winning bidder** is the bidder with the highest bid and only that bidder makes a payment to the auctioneer. The amount it pays to the auctioneer is called the auction's **clearing price**. Many auctions incorporate a **floor price** (also known as a **reserve price**), which is a price below which the seller is unwilling to sell. In an auction with a floor price, only bids that exceed the floor price are considered.

52. Auctions are popular when different bidders might value an item differently because they allow bidders' information about different goods to be reflected in prices, a process called

[63] Other types of auctions may elicit information from the same bidder *multiple* times during the auction: for example, in "open outcry" English auctions (an auction format commonly used for art auctions at, say, Sotheby's), bidders may make multiple sequential bids for an item, each of which must be higher than the current highest bid, and the bidder who makes the final bid wins the auction at a price equal to its bid.

price discovery.[64] Sealed-bid auctions are particularly suitable for online display ad impressions because of their ability to discover prices and allocate items rapidly to bidders whose values may be quite different and changing over time.

2. How Auction Designs Affect Publisher and Advertiser Incentives

53. In much of this report, I assume that an advertiser's objective in an auction is its **advertiser surplus** (also known as **advertiser profit**), which is equal to the difference between its value for the impression if it wins it and the price it pays.[65] An advertiser chooses a bid to maximize the *expected* value of this advertiser surplus, given its prediction of the behavior of other bidders. A publisher's objective is its **revenue**, which is equal to the payment it receives for an impression. The publisher chooses its floor price in the auction, which is the minimum payment it is willing to accept for the impression (discussed further in Section III.C.3.d) to maximize its *expected* revenues, given its predictions of the behavior of bidders.[66] If each agent chooses actions to maximize its benefits of participating in the auction, assuming that all other agents also maximize their returns, and if agents' forecasts of others' choices are statistically accurate, then their

---

[64] *See, e.g.*, Milgrom, P. (2017). *Discovering prices: auction design in markets with complex constraints*. Columbia University Press, at 46 ("Even when prices to guide efficient resource allocation exist in theory, the practical problem of finding those prices can still be daunting. […] The best way to find such prices is often an auction of some kind.").

[65] Advertisers may have other campaign objectives, but as I discuss in Section VII.C, *any* rational strategy for an advertiser (*i.e.*, any strategy that purchases a set of impressions in the *least-cost* way) leads to choosing bids for each individual impression that maximize expected advertiser surplus given some value for the impression.

[66] Publishers and advertisers may make decisions other than floor prices and bids, including where to sell or bid for an impression and how many impressions to offer. Where these other decisions are relevant to my analysis, I also assume that publishers and advertisers make these decisions in line with their incentives.

choices are said to be in **(Bayes-Nash) equilibrium**.[67] Equilibrium analysis is a common basis of theoretical predictions about auctions because it accounts for the incentives of all participants.

54. In much of this report, I adopt the common approach of assuming that advertisers and publishers respond to the rules and conditions they face to maximize their individual returns from online display advertising. There are several reasons that I believe this to be the relevant standard for this case.

55. *First*, there are vast sums of money at stake for many of the publishers and advertisers in the industry. Many also have large teams of engineers, economists, and marketing experts devoted to maximizing returns by finding all possible improvements in advertising yields. The largest advertisers and publishers also make up a significant fraction of the online display advertising industry, suggesting that the strategic sophistication of these agents is especially relevant to auction outcomes.[68] Academic research has found that bidders with large stakes in auctions typically behave in line with their incentives, as predicted by equilibrium analysis.[69]

[67] Bayes-Nash equilibrium is a standard solution concept of game theory. *See* Harsanyi, J. C. (1967). Games with incomplete information played by "Bayesian" players, I-III. Part I. The basic model. *Management Science*, *14*(3), 159-182; Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT Press, at 3, 11-14.

[68] For example, in 2022, more than [redacted]% of Google Ads' US web spending across AdX, AdSense, and non-Google exchanges was by advertisers spending more than $100,000 on the platform. For the same criteria, more than [redacted]% of DV360 spending was by advertisers spending more than $100,000. *See* GOOG-AT-MDL-DATA-000486626 to -8277 (Google Ads); GOOG-AT-MDL-DATA-000561263 to -420 (XBridge DV360). This result was generated using [redacted] in my supporting materials, and the output is saved in [redacted]

[69] *See, e.g.*, Doraszelski, U., Lewis, G., & Pakes, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, *108*(3), 565-615; Hortaçsu, A., & Puller, S. L. (2008). Understanding strategic bidding in multi-unit auctions: A case study of the Texas electricity spot market. *The RAND Journal of Economics*, *39*(1), 86-114; Hortaçsu, A., Luco, F., Puller, S. L., & Zhu, D. (2019). Does strategic ability affect efficiency? Evidence from electricity markets. *American Economic Review*, *109*(12), 4302-4342.

56. *Second*, publishers and bidders in online display advertising typically participate in many auctions, with this repeated interaction creating opportunities to experiment, learn, and improve strategies over time. Empirical evidence from online display advertising auctions suggests that agents learn to respond to auction design changes over time, and eventually come to adopt near-profit-maximizing strategies.[70] This research suggests that strategic adaptation is not always immediate and that there is heterogeneity in the speed of learning, which implies that evidence about the impact of new programs gathered over short periods of experimentation must be evaluated with care: it may fail to capture eventual strategic adaptations and heterogeneity in effects over time and across agents.

57. *Third*, intermediaries that serve publishers and advertisers compete for business by offering optimization tools, so that even smaller advertisers and publishers have access to sophisticated tools to optimize performance. The *stated* objective of many tools therefore offers a window into what they believe their customers—publishers and advertisers—may want. These statements clearly support the hypothesis that publishers and advertisers seek to optimize returns.

58. Here are some examples. On the supply-side, Microsoft's Xandr offers "yield optimization tools to maximize the value of […] inventory."[71] Pubmatic offers tools to "maximize advertising revenue."[72] Criteo's Commerce Grid SSP gives "media owners the

[70] *See, e.g.*, Goke, S., Weintraub, G. Y., Mastromonaco, R., & Seljan, S. (2022). Bidders' responses to auction format change in internet display advertising auctions. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4021221.

[71] Xandr, "Publisher Platforms," Microsoft Advertising (accessed Oct. 25, 2023), https://about.ads.microsoft.com/en-us/solutions/xandr/publisher-platforms-scaled-buying-selling-solutions.

[72] Pubmatic, "Pubmatic SSP: Maximize Advertising Revenue and Control How Your Audiences are Accessed," Pubmatic (accessed Oct. 25, 2023), https://pubmatic.com/products/pubmatic-ssp-for-publishers/.

control to optimize the monetization of their inventory."[73] Google Ad Manager offers "yield management solutions" including to "[m]aximize revenue in the unified first price auction."[74] Clickio offers floor price optimization tools and other services for publishers to "guarantee maximum revenues."[75] On the demand-side, The Trade Desk offers tools to advertisers to "[m]aximize [their] digital investment" by "[r]each[ing] the right audience and spend[ing] [their] budget more efficiently."[76] Xandr's DSP offers optimization tools to allow advertisers "to get the greatest benefit for the money spent on campaigns."[77] Criteo's smart bidding tool for its advertiser customers "ensures a bid placed for any display opportunity offers the highest yield for clients and publishers."[78] Google Ads offers "[t]ools to help [an advertiser] optimize […] bids" and help with "maximizing [its] Google Ads budget."[79] Google's Display & Video 360 (DV360) offers automation tools

73 Criteo S.A., Form 10-K, United States Securities and Exchange Commission (2022), https://www.sec.gov/Archives/edgar/data/1576427/000157642723000023/crto-20221231.htm.

74 Google, "Get comprehensive yield management with Google Ad Manager," Google Ad Manager (accessed Oct. 25, 2023), https://admanager.google.com/home/resources/feature-brief-yield-management/.

75 Clickio, "Monetization" (accessed Jul. 7, 2024), https://clickio.com/monetization/.

76 The Trade Desk, "DFP Features[:] Integrated planning and activation (accessed Oct. 31, 2023), https://www.thetradedesk.com/us/our-platform/dsp-demand-side-platform/plan-campaigns.

77 Xandr, "Understanding Optimization," Xandr Documentation Center (Feb. 7, 2022), https://docs.xandr.com/bundle/monetize_monetize-standard/page/topics/understanding-optimization.html.

78 Criteo, "What You Need to Know About First-Price Auctions and Criteo," Criteo Updates (Sep. 26, 2023), https://www.criteo.com/blog/first-price-auctions/.

79 Google, "Evaluate and optimize your bids," Google Ads Help (accessed Oct. 31, 2023), https://support.google.com/google-ads/answer/7085711?hl=en; Google, "10 tips for Google Ads budget management," Google Ads Resources (Mar. 20, 2023), https://ads.google.com/intl/en_us/home/resources/articles/stretching-your-google-ads-budget/. *See also* Deposition of [redacted] at 370:6-8 (Sep. 17, 2021), GOOG-AT-MDL-007173084, at -454 ("The goal of predicted highest other bid is to optimize towards the advertiser's surplus."); Design Doc, "AdX/AdMob first price bidder - for perf" (Sep. 26, 2019), GOOG-DOJ-AT-00573492, at -495 ("Surplus maximization to determine the optimal bid: First-price bid is the one that maximizes the advertiser surplus = advertiser value - payout.").

to "help make optimal bids to help improve campaign performance."[80, 81] The stated objectives of these tools offered by intermediaries suggest that optimizing returns is a leading goal of advertisers and publishers and that their customers outsource many decisions to intermediaries to reduce the burden of understanding all the details of auction rules and optimizations themselves.

3. Types of Sealed-Bid Auctions

59. Auctions can use a variety of rules to determine allocations and payments. Different auction designs create different incentives for auction participants.

a) Second-Price Auctions and Threshold Pricing

60. A common auction practice early in the history of the online ad industry was the use of **second-price auctions**, in which the highest bidder for an impression wins and pays a price equal to the larger of the floor price or the second-highest bid. The second-price auction rule is the sealed-bid implementation of the well-known **ascending auction**,[82] in which the auctioneer asks for bids at the floor price and gradually increases the price until just one bidder remains. The winning bidder in an ascending auction pays a price that is determined by the drop-out price of the *second-to-last-remaining* bidder, just as the winning bidder in the second-price sealed-bid auction pays a price determined by the bid of the *second-highest* bidder.

[80] Google, "Enhanced automation," Display & Video 360 Help (accessed Oct. 31, 2023), https://support.google.com/displayvideo/answer/6130826?hl=en.

[81] The surplus-maximization objective of DV360 is also clearly captured in the design of Poirot, as discussed further in Section VII.

[82] *See* Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37.

61. The second-price auction is the canonical member of a class of auctions that uses **threshold pricing**, meaning that the price paid by the winning bidder does not depend on its winning bid but is instead the lowest amount (called the **threshold**) that the winning bidder could have bid to win the auction, holding all other bids fixed. Auctions with threshold pricing can vary in their winner-selection rules: the second-price auction is the auction using threshold pricing in which the highest bid always wins.[83] All auctions with threshold pricing have the following important property: once an advertiser has determined its value for an impression—the maximum price it is willing to pay for that impression—it can maximize its profit in the auction simply by bidding that value.[84, 85] An auction with this property is called **bidder-truthful** (or sometimes **incentive-compatible** or **strategy-proof** or **truthful**).

62. To see why it is always optimal with threshold pricing for the advertiser to bid its value, v, there are two cases to consider. First, suppose that the minimum bid needed to win the auction is some amount x, which is less than v. Then, bidding the advertiser's value v will win the auction and, by the threshold pricing rule, the winner will pay x. Any other bid either wins the auction at the same price or loses the auction (guaranteeing zero payoffs), so no other bid can do better: bidding v maximizes the bidder's surplus. Second, suppose

[83] As another example, a publisher's online ad auction might specify that some favored bidder wins an impression if its bid plus 25% is higher than any other bid; otherwise, the item goes to the highest other bidder. Using this rule to determine the winner results in the favored bidder, when it wins, paying a discounted threshold price. For example, if the highest bid from a regular customer is 100 and the favored bidder bids more than its threshold of 80, then it wins and pays a price of 80, even if its own bid was, say, 90.

[84] Milgrom, P., & Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, *128*(1), 1-31, at 16-18 ("Proposition 3. Any threshold auction is strategy-proof. Conversely, any strategy-proof direct auction has a monotonic allocation rule, and if V=R^N_+, it must be a threshold auction.").

[85] Bidder-truthfulness applies to a single auction for an impression. When the same agents interact in multiple auctions, advertisers may have different incentives leading them not to bid their values. For example, an advertiser may not want to reveal that it has a high value for an impression if it anticipates that the publisher will change its floor price in subsequent auctions to take advantage of that information.

that the minimum bid to win the auction is some amount y, which is greater than v. Then, bidding the advertiser's value v will lose the auction, but any winning bid would require that the advertiser pay a price y, which is more than its value. So, in this second case, too, there is no bid that earns the advertiser more than bidding its value v.

63. Bidder-truthful auctions reduce bidding errors and the costs of bidding because they eliminate any need for an advertiser to assess who else might be bidding, how much they might bid, or the publisher's floor price. In non-bidder-truthful auctions, each advertiser's bid depends on all of these factors. I have previously advised auctioneers to adopt bidder-truthful auctions, highlighting the importance of easy bidding.[86]

64. Other auctions using threshold pricing besides the second-price auction were used in several of Google's programs, as described later in this report. Threshold pricing is special not just because it results in bidder-truthful auctions, but also because these are the *only* bidder-truthful auctions. There are no others.[87] Any auction in which a winning bidder pays something other than its threshold price is not bidder-truthful, and as a consequence, incentivizes bidders to choose bids that differ from their values for the good being sold.

[86] When the US Federal Communications Commission sought to repurchase certain television broadcast rights, I advised a bidder-truthful auction so that for any television broadcaster, "the hardest part of bidding will be to determine its value of continuing to broadcast. Once it knows that value, the rest is easy. The bidder cannot do better than to agree to accept any price greater than its value of continuing to broadcast and then to exit if its offered price falls lower than that. By bidding in this way, the station will obtain its best possible price[…]" Auctionomics and Power Auctions, "Incentive Auction Rules Option and Discussion," FCC (Sep. 12, 2012), https://docs.fcc.gov/public/attachments/FCC-12-118A2.pdf, at 3.

[87] For proof, notice that if the price paid by a winning bidder depends on its bid, then instead of bidding truthfully, a bidder prefers to make the winning bid that results in the lowest price. *See* Milgrom, P., & Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, *128*(1), 1-31, at 16-18.

65. Google recognized the advantages of bidder-truthful auctions, explaining them as follows: "It's faster, less costly, and more fair to the less sophisticated advertisers to structure the auction in favor of true value."[88] The lower transaction costs associated with bidding in a bidder-truthful auction encourage advertisers to participate on Google's platform, which increases thickness, tending to improve the efficiency of its allocations and increase the prices paid to publishers.[89]

66. Second-price auctions have an additional important benefit: if bidders bid truthfully, then the price determined by the auction for each impression is a **market-clearing price**. A market-clearing price is one at which supply equals demand—one bidder is willing to pay the price and no losing bidder is willing to pay more. Whenever an impression is transacted at a market-clearing price, that transaction maximizes the economic welfare that can be created from the sale of the impression. It can be proved mathematically that the second-price auction is the only bidder-truthful auction that always transacts the impression at a market-clearing price.[90]

b) First-Price Auctions

67. In a first-price auction, the bidder with the highest bid wins and pays a price equal to its bid for the item (unless the highest bid is below the floor price, in which case the item is unallocated). In contrast to the second-price auction, *every* bidder in a first-price auction needs to **shade** its bid—that is, choose a bid *less* than its value—in order to stand any

[88] "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -085.

[89] *See, e.g.*, Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37.

[90] *See* Milgrom, P. (2004). *Putting auction theory to work*. Cambridge University Press, at 71-73.

chance of making a profit.[91] From a bidder's perspective, its optimal bid shading results from a tradeoff: a higher bid increases the probability of winning an impression, but it also increases the bidder's price for the impression if it wins. The optimal shading calculation depends on the bidder's estimates or guesses about what others might bid: the bidder should reduce its bid more if it expects lower bids from others. Guessing the identities and bids of others for each different impression is a costly and error-prone activity that can lead to inefficiency when bidders' guesses are wrong. In a first-price auction, the winner may be a bidder who does not have the highest value for the impression. Even when the impression is allocated correctly, the winning bid may not be a market-clearing price: after seeing the outcome, some losing bidders may have been willing to pay more than the winning bid for that impression.

68. While second-price auctions are bidder-truthful, first-price auctions have their own advantages. One important benefit of the first-price auction is its transparency. Unlike in a second-price auction, the winning bidder always pays its bid, so it does not need to trust the auctioneer's computations or reports of floor prices and other bids in order to confirm that the price it pays was determined properly. Among sealed-bid auctions, only first-price auctions are what auction theorists have called **credible auctions**, which means that the auctioneer has no profitable way to cheat that is not immediately detectable by the winning bidder.[92] This credibility property is one reason that first-price

[91] Bid "shading" is a term of art in the economic theory of auctions, dating back to at least the seminal work of William Vickrey. *See* Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37, at 12-13 ("On the other hand, if traders have a fairly confidently held expectation that the equilibrium price will fall within a certain narrow range, there may be an indirect community of interest in shading the reported demand and supply curves outside this range in the direction of greater inelasticity"). The term has also been used to describe bid optimization in the online display advertising industry.

[92] Akbarpour, M., & Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, *88*(2), 425-67, at 427.

auctions have been preferred in some applications. I discuss additional advantages of the first-price auction format for online display advertising in Section III.C.4 below.

c) Other Auction Formats

69. Other auction formats exist and lead to different incentives for auction participants. One example is the **1.5-price** auction, which charges the winning bidder an amount halfway between its own bid and the second-highest bid. In the 1.5-price auction format, bidders optimize by shading their bids, but to a lesser extent than in the first-price auction. Another example, relevant in online display advertising, is a non-transparent auction, in which the auctioneer might claim to calculate winners and payments according to one rule, but actually charges bidders according to another rule. Because it is so easy to detect the first-price rule, an auctioneer running a non-transparent auction might claim to use a second-price auction but actually use the 1.5-price rule, hoping to confuse bidders into bidding too much, increasing the auctioneer's profit. When non-transparent auctions are possible, optimizing bidders must rely on data about past auction performance and/or on experiments to determine optimal bids into the auction.

d) How Publishers Set Floor Prices

70. Setting a floor price can help increase publisher revenues by ensuring that no impression is ever sold at a very low price. Determining the optimal floor price, however, requires a subtle strategic computation. Raising the floor price may sometimes increase auction revenue, but if it is set too high, there might be no bidder willing to bid more than the floor price. Then, the impression may remain unsold, leading to zero revenues for the publisher and zero surplus for the advertiser.

71. To determine a floor price that maximizes the expected revenue on an impression, a publisher must make a probabilistic assessment of the costs and benefits of each possible floor price, trading off the possible benefit of an increased price in the case that the impression sells against the cost of an increased likelihood that the impression is unsold.

72. As an example of this computation, suppose that a publisher is selling an impression to a single bidder using a second-price auction and knows that it is worth either $1, $2, or $3 to that bidder, each with equal probability. If the publisher sets a floor price of $3, it expects to sell the impression for $3 with probability ⅓, yielding an expected revenue of $1. If it sets a floor price of $1, it will always sell the impression for $1. To maximize its expected revenue, the publisher should set a floor price of $2, which sells the impression with probability ⅔, leading to an expected revenue of $1.33.[93]

73. The above example shows that setting a floor price properly can increase revenues for a publisher in a second-price auction. Setting a floor price can also increase revenues for a publisher in a first-price auction, by reducing the extent to which some bidders shade their bids.

74. In practice, in repeated auction settings, publishers often use simulations on recent auction data or experiments (sometimes known as "A/B tests") of different floor prices on live auctions, and choose the floor price that led to the highest revenue in those simulations or experiments.[94] [redacted]

[93] The publisher could choose floor prices other than $1, $2 or $3, but—by similar calculations to the ones above—any such floor price leads to lower expected revenues than the floor price of $2.

[94] *See, e.g.*, Rhuggenaath, J., Akcay, A., Zhang, Y., & Kaymak, U. (2022). Setting reserve prices in second-price auctions with unobserved bids. *INFORMS Journal on Computing*, *34*(6), 2950-2967.

[redacted][95] [redacted]"[96] There are supply-side intermediaries who assist in this process, as well, including Clickio, which offers floor price optimization tools and other services for publishers to "guarantee maximum revenues."[97]

4. The Industry's Switch from Second-Price Auctions to First-Price Auctions

75. Early online ad exchanges used second-price auctions to allocate online display advertising impressions.[98] But as online display advertising evolved, new challenges emerged that threatened to undermine the performance of second-price auctions and eventually led many ad exchanges to switch auction rules.

76. *First*, a DSP bidding on behalf of multiple advertisers could increase its profits by submitting only one bid into a second-price auction, instead of submitting bids on behalf of all its advertisers,[99] while continuing to charge advertisers the same threshold prices (preserving the bidder-truthfulness of the auction). If a DSP pursues that strategy, then whenever it hosts the two highest bids for an impression, it ends up paying less for the impression than it would if it submitted bids on behalf of all its advertisers. For example, suppose that the DSP's two highest bids are $12 and $10 and the highest bid submitted by

---

[95] [redacted]

[96] [redacted]

[97] Clickio, "Monetization" (accessed Jul. 7, 2024), https://clickio.com/monetization/.

[98] Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -267 (Figure: "Impression Share (US Ad Inventory)"; "Fair Second-Price Auction": 75% in December 2017, 33% in March 2018).

[99] Submitting bids determined by the values of the two highest-value advertisers would lead to the same outcomes in a second-price auction as submitting bids on behalf of all advertisers.

any other demand source is $8, all higher than the auction's floor price. If the DSP submits both of its bids to the auction, the DSP wins the auction and pays its second bid of $10. If the DSP instead submits just the bid of $12 (and assuming no change in the auction's floor price and the bids of other bidders), then it wins the auction and pays the new second-highest bid of $8. This does not imply, however, that the winning advertiser benefits: by charging the winning advertiser its threshold price, that advertiser pays $10 and the DSP pockets the $2 difference between that amount and the auction's clearing price ($10 - $8), increasing its profits.[100] As in this example, the general result of this DSP's **one-bid strategy** is that the publisher loses revenue. With this strategy, the price paid to the publisher is not a market-clearing price (because at least two advertisers would be willing to pay more than the price of $8), and the DSP can earn an additional profit just by keeping the price difference. Some non-Google DSPs are reported to have done exactly that,[101] and Google observed that many bidders did not submit two bids into the AdX auction.[102]

77. *Second*, some publishers adopted a tactic called **multi-calling**.[103] Instead of calling the exchange with the floor price that it would set in a single auction, these publishers would first call the exchange with an inflated floor price. If advertisers bid truthfully, this first

[100] To see that bidders' incentives are not affected by the DSP strategy, note that the allocations and payments for the bidders are the same under the DSP strategy as in a second-price auction.

[101] *See, e.g.*, Ross Benes, "Ad buyer, beware: How DSPs sometimes play fast and loose," Digiday (May 25, 2017), https://digiday.com/marketing/dsp-squeeze-buyers/ ("[S]ome [DSPs] bill based on a clearing price for the auction that occurred within the DSP's platform that can be higher than the price that won the impression on the open exchange, and the DSP will keep the difference.").

[102] Presentation, "Understanding the AdX Auction" (Oct. 2014), GOOG-DOJ-12443562, at -567 ("[O]ther bidders are allowed to [second-price themselves], but often do not").

[103] "Solving the Multi-Call problem" (Nov. 25, 2019), GOOG-AT-MDL-001397473, at -475 ("Background 2: The multi-call problem").

call might cause a bidder to win the auction at the inflated price. If there was no winner at the high price, the publisher would call the exchange again, offering the same impression with a lower floor price. As a result, if the publisher engages in multi-calling, it is not generally optimal for an advertiser to bid its value: the winning advertiser may be able to buy the impression at a lower price by reducing its bid, letting the impression go unsold in the first auction, and then winning in a later auction at a lower price. Multi-calling destroys the simplicity of threshold pricing and so complicates bidding for advertisers, forcing them to strategize about how best to respond to the publisher's practice and make guesses about the publisher's true floor price and about others' bids. Multi-calling also increases processing costs and adds latency, damaging the end user's online experience and leading to a reduction in advertising effectiveness.[104]

78. *Third*, as online display advertising platforms evolved to accept more bids from different demand sources in each auction, **self-competition**—which occurs when an advertiser submits multiple bids for the same impression—became a larger concern for advertisers. Self-competition can occur as a result of **advertiser multi-homing**,[105] in which an advertiser uses multiple DSPs to submit bids for an impression, or **DSP multi-homing**,[106]

[104] *See, e.g.*, Oded Poncz, "Traffic duplication might be a bigger problem than ad fraud," AdExchanger (Jan. 11, 2016), https://www.adexchanger.com/data-driven-thinking/traffic-duplication-might-even-be-a-bigger-problem-than-ad-fraud/ ("Another side effect of bid request duplication is that re-auctioning a bid takes time. In some cases, this could even become apparent to the end user.").

[105] In a 2021 survey, respondent advertisers and ad agencies (who all spent a minimum of $1M annually on digital ads) used an average of [redacted] DSPs and planned to use [redacted] DSPs the following year. *See* Advertiser Perceptions, "DSP Report: Demand-Side Platforms" (2021), GOOG-DOJ-AT-02524665, at -666, -670. *See also, e.g.*, [redacted]

[106] *See, e.g.*, [redacted]

in which a DSP submits bids into multiple exchanges on behalf of a single advertiser. When the same bidder submits multiple bids in a second-price auction, it may wind up making both of the two highest bids, with its own second-highest bid setting its price. In such cases, the advertiser would pay a lower price if its second bid were lower or omitted. Bidders would need to adjust their bids to avoid this possibility, requiring more complicated bidding strategies.

79. *Fourth*, some exchanges tried to increase their profits by using **non-transparent auctions**, claiming to calculate the winner's price using a second-price rule, but actually charging winners a larger amount, for example by using the 1.5-price rule described earlier.[107] Such practices make participation in the auction less safe for advertisers and reduce trust in online display advertising intermediation more generally, harming other exchanges and economic welfare. In order to protect against non-transparent auctions, bidders would need to invest in costly technology to detect exchanges using non-second-price auction rules and to optimize bids into those exchanges.

80. Many intermediaries adopted practices to reduce the potential harms associated with the four challenges I described above,[108] and eventually ad exchanges across the industry responded by switching away from a second-price auction format. Most exchanges,

[107] Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -635 ("Dirty [auctions are] called second price, but really more like first price"; "Dirtiness is introduced using a new type of floor called a'Soft-floor' […] to achieve a continuum of auctions from second price […] to 1st price […], opaque to the advertiser"; "All these auctions [on exchanges United, AdX, AppNexus, OpenX, Pubmatic] are 'second price,' however the auction discount (cost/bid) varies widely.").

[108] For example, Projects Bell and Poirot at Google were designed to respond to multi-calling and non-transparent non-second-price auctions, respectively. Non-Google display advertising intermediaries also introduced similar products, as demonstrated in Table 1 and discussed below.

including Google's AdX, made the switch to a first-price format between 2017 and 2019.[109]

81. The transition to a first-price auction reduced or eliminated the four concerns I listed above. *First*, in a first-price auction, a DSP cannot reduce a publisher's revenue by suppressing its lower bids because each bidder always pays its highest bid. *Second*, bidders that optimize their bids in first-price auctions using experimentation can reverse some of the losses to multi-calling, making it less likely for that tactic to be profitable.[110] *Third*, in a first-price auction, a bidder gains some protection against self-competition, because unlike in a second-price auction, its losing bids do not affect its own price.[111] *Fourth,* no auctioneer can profit from deceiving a bidder about the correct price: the winning bidder can easily check whether its payment equals its bid.

[109] *See* Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -267 (Figure: "Impression Share (US Ad Inventory)"; "First-price auction": [redacted]% in December 2017, [redacted]% in March 2018). These figures would not include AdX's later transition to a first-price auction in 2019. *See also* Email from [redacted]., "Re: Offering 1st price to publishers?" (Sep. 5, 2017), GOOG-DOJ-05272070, at -075 ("Other exchanges started the migration to 1st price auction recently arguing th[at] it is the best way to integrate it with HB.").

[110] In a first-price auction, bidders typically experiment to identify optimal bids, and—in the face of multi-calling tactics by publishers—such experiments would identify benefits from reducing bids more to multi-calling publishers than non-multi-calling publishers, resulting in similar effects as Project Bell, discussed in Section V below.

[111] A multi-homing advertiser still needs to be mindful of self-competition in first-price auctions, in case one of its buying tools learns to bid more aggressively in response to that same advertiser's higher bids via different buying tools (driving up prices for the advertiser). This is one reason for the emergence of so-called "supply path optimization tools" which help bidders optimize the platforms on which they bid for inventory. *See* Yuyu Chen, "WTF is supply path optimization?," Digiday (May 22, 2023), https://digiday.com/media/what-is-supply-path-optimization/("[Supply path optimization] is essentially an algorithm used by demand-side platforms to streamline how they interact with supply-side platforms. Each DSP has developed its own strategy for supply path optimization: Some use it to pick up the bids that are most relevant and have the highest chance of winning, while others use it to turn off SSPs that are not implementing second-price auctions, according to Tom Kershaw, CTO for ad exchange Rubicon Project. […] Two major reasons [why it is important for DSPs]: Bid duplication and various auction mechanisms used by SSPs.").

D. Google's Products

82. Along with many other companies, Google provides a variety of intermediary services in online display advertising. In this report, I study the practices of a number of Google's intermediaries: on the demand-side, **Google Ads** and **Display & Video 360 (DV360)**, and on the supply-side, **Google Ad Manager (GAM)**, which incorporates advertising exchange functionality (formerly known as **AdX**) and publisher ad server functionality (formerly known as **DFP**).

1. Display & Video 360

a) How Advertisers Use DV360

83. **DV360** (formerly known as DoubleClick Bid Manager or DBM) is a demand-side platform that offers tools for "planning [display advertising] campaigns, […] designing and managing creative[s], organizing audience data, finding and buying inventory, and measuring and optimizing campaigns,"[112] enabling advertisers to "manage their reservation, programmatic, and programmatic guaranteed campaigns across display, video, TV, audio, and other channels, all in one place."[113] As I described above, DV360 offers automated tools to "help make optimal bids to help improve campaign

[112] Google, "Display & Video 360 overview," Display & Video 360 Help (accessed Jan. 21, 2024), https://support.google.com/displayvideo/answer/9059464?hl=en.

[113] Google, "Display & Video 360: An integrated solution for end-to-end advertising campaigns," Display & Video 360 (accessed Oct. 31, 2023), https://services.google.com/fh/files/misc/display_and_video_360_product_overview.pdf.

performance."[114] DV360 allows advertisers to purchase ads from many sell-side platforms, including AdX, Index Exchange, OpenX, Rubicon, and others.[115]

84. Advertisers on DV360 can set up their campaigns in multiple ways. In the early days of DV360, most advertisers used DV360 to set up **fixed CPM** campaigns, in which the advertiser would report to DV360 the characteristics of the impressions they would like to purchase with each campaign and a fixed CPM—cost per mille (thousand impressions)—and DV360 would use that CPM to bid in auctions for those impressions.[116] Fixed CPM bidding is still offered as an option on DV360, but between 2017 and 2020, the majority of DV360 advertisers switched to **automated bidding** campaigns,[117] in which the advertiser reports an objective to DV360 (for example, "maximize clicks" or "maximize conversions" subject to a budget) and DV360 applies prediction and optimization algorithms to "dynamically determine the optimal bid price for a given impression for an advertiser."[118] More recently, DV360 introduced **custom**

[114] Google, "Enhanced automation," Display & Video 360 Help (accessed Oct. 31, 2023), https://support.google.com/displayvideo/answer/6130826?hl=en.

[115] Sissie Hsiao, "How our display buying platforms share revenue with publishers," Google Ad Manager (Jun. 23, 2020), https://blog.google/products/admanager/display-buying-share-revenue-publishers/ ("Using Display & Video 360, these advertisers can buy ads on more than 80 publisher or sell-side platforms including AT&T, Comcast, Index Exchange, OpenX, Rubicon Project, MoPub and others.").

[116] Presentation, "DV360 optimizations ENG deep dive" (Jan. 24, 2020), GOOG-DOJ-11733552, at -553 ("DV360 three years ago: Mostly fixed CPM manual bidding").

[117] Presentation, "DV360 optimizations ENG deep dive" (Jan. 24, 2020), GOOG-DOJ-11733552, at -555 ("Impact on auto-bidding adoption: [redacted]% now").

[118] Comm Doc, "Optimized Fixed Bidding in DV360" (Mar. 2019), GOOG-DOJ-05326023, at -023.

bidding, which allows advertisers to input their own algorithms to determine bids on an impression-by-impression basis.[119]

b) How DV360 Determines Bids and Payments

85. When DV360 receives a bid request from an exchange, it first identifies the advertiser campaigns that are targeting impressions with the characteristics associated with the bid request.[120] As I discuss further in connection to Project Elmo in Section VI below, in order to ensure that advertiser budgets are not depleted too quickly, DV360 then applies its **budget throttling** algorithm to determine a selection of eligible advertisers for participation in the auction for the impression. Among those advertisers, DV360 then determines the advertisers with the highest bids for the impression net of any fees and submits bids to the exchange.[121] If DV360 expects that the exchange is using a non-second-price auction, the submitted bids are adjusted using a bid optimization

[119] Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -270 ("Rule based custom bidding is a new bidding strategy that we are building together with the KIR team. It allows our advertisers and agencies to express what they really value in advertising and optimize towards that value. […] Once they know what features and signals they need, they can […] [w]rite a script on how to evaluate the value of a query in the most generic form. Upload the script to us. We take this script and score all of their historical impression data. Then we learn from that data to bid optimally on future received queries."); Google, "Custom bidding overview and limitations," Display & Video 360 Help (accessed Oct. 31, 2023), https://support.google.com/displayvideo/answer/9723477?hl=en("Custom bidding scripts and goals let you define the value of an impression that aligns with your campaign's goals. Display & Video 360 uses the algorithm from your custom bidding scripts or goals to determine the score it assigns to impressions. The outcome of previously scored impressions helps train the custom bidding model to help optimize your bidding strategy.").

[120] Presentation, "DV360 product and architecture" (Oct. 31, 2018), GOOG-DOJ-15407321, at -363 ("DBM finds fitting Campaigns").

[121] Presentation, "DV360 product and architecture" (Oct. 31, 2018), GOOG-DOJ-15407321, at -363 ("Mixers run auctions between campaigns/systems").

algorithm called Poirot that makes bidding simpler for advertisers.[122] I discuss Poirot in detail in Section VII.

86. If the DV360 bid wins the auction, the advertiser with the highest bid in DV360 is allocated the impression and pays the clearing price of the exchange auction, plus any fees charged by DV360. The fees charged to a given advertiser are determined by contracts negotiated between Google and the advertiser, and they include platform fees that are typically a fixed percentage of the total price of each impression.[123] Prior to August 2023, DV360 advertisers could also elect to be charged for impressions on a "pay per outcome" basis, in which they paid only if the end user engaged with the ad according to some predetermined outcome measure (similar to payments by some Google Ads advertisers, discussed below).[124, 125]

[122] Presentation, "DV360 product and architecture" (Oct. 31, 2018), GOOG-DOJ-15407321, at -363 ("DBM Frontend translates the Bid for the Exchange"); Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -809 ("Over [redacted] of DBM bidding goes through third-party exchanges […] Fixed CPM bidders have the same bid in these unclean exchanges as they do in clean exchanges […] The goal of Poirot is to discover the exchanges that deviate from second pricing and bid appropriately on these").

[123] *See* "FY22 - DV360 (DBM) Narrative" (Jul. 12, 2023), GOOG-AT-MDL-008930706, at -713 to -715 ("DBM Partners (a term used interchangeably with advertisers or customers) must have an Advertising Platform Agreement (APA) or Affiliates Adopting Agreement (AAA) in place with Google in order to use DoubleClick products. […] The standard DBM deal structure is to charge customers a platform fee (*i.e.* license/technology fee) based on transaction type."). Google may also charge advertisers on DV360 ad serving fees for third-party services or technologies, and fees for access to certain types of data from third-party providers.

[124] *See* Design Doc, "Pay per Outcome in DBM" (Feb. 10, 2018), GOOG-AT-MDL-009590288, at -289 ("In this document, we discuss a detailed design to make outcome based buying possible on DBM.").

[125] This service has now been deprecated. *See* Google, "Coming soon: April 17, 2023 edition," Display & Video 360 announcements (Apr. 17, 2023), https://support.google.com/displayvideo/answer/13511229 ("Outcome based buying will be deprecated on August 1 (previously communicated as July 1), 2023").

2. Google Ads

a) How Advertisers Use Google Ads

87. **Google Ads** (formerly known as AdWords) is a buy-side tool that permits advertisers to create ad campaigns that run across different formats, including search and display ads.[126] In this report, I focus on the online display advertising functions of Google Ads on web properties other than those owned and operated by Google.

88. Advertisers on Google Ads specify campaign goals (*e.g.*, maximizing clicks or conversions) and constraints for a campaign (*e.g.*, types of end users to target, maximum bids or budgets), and Google Ads uses that information to determine bids for impressions.[127] Advertisers can also specify "manual" bids for impressions on a CPC (cost-per-click) or CPM basis.[128] In this report, for simplicity, I use the word "bid" to refer to either the manual bid reported by the Google Ads advertiser or the bid determined by Google Ads as a function of the advertiser's reported campaign goals, whichever is relevant for that advertiser.

89. Originally, beginning in 2003, advertisers could use Google Ads to buy third-party display advertising inventory only from publishers using AdSense, a Google supply-side

[126] Google, "Google Ads" (accessed Jan. 5, 2024), https://ads.google.com/intl/en_us/home/.

[127] Google, "Determine a bid strategy based on your goals," Google Ads Help (accessed Oct. 31, 2023), https://support.google.com/google-ads/answer/2472725 ("Depending on which networks your campaign is targeting, and whether you want to focus on getting clicks, impressions, conversions, or views you can determine which [bidding] strategy is best for you.").

[128] Google, "Determine a bid strategy based on your goals," Google Ads Help (accessed Oct. 31, 2023), https://support.google.com/google-ads/answer/2472725 ("Manual CPC bidding: This lets you manage your maximum CPC bids yourself. […] CPM: With this bid strategy, you'll pay based on the number of impressions (times your ads are shown) that you receive").

product.[129] After the acquisition of DoubleClick and the 2009 launch of AdX 2.0, advertisers could also use Google Ads to purchase inventory from AdX.[130] Since the beta launch of **AwBid** in May 2013 and its general launch in June 2015, advertisers have also been able to use Google Ads to purchase some types of inventory through non-Google ad exchanges.[131] Google Ads has its own bid optimization program for non-Google exchanges called **Marple**, which worked similarly to Poirot on DV360, determining optimal bids for Google Ads advertisers in non-second price auctions.[132] The AwBid program has grown over time, but the majority of Google Ads' total online display advertising spend is on AdX.[133]

b) How Google Ads Determines Bids into AdX: The Internal Auction

90. When Google Ads is called to bid on an impression, it uses an internal process which I will call the **Google Ads internal auction** (and which Google employees refer to as the

[129] Email from [redacted] to [redacted], "Final Google AdSense press release" (Jun. 8, 2003), GOOG-DOJ-01805208, at -208 ("With Google AdSense, publishers serve text-based Google AdWords ads on their site and Google pays them for clicks on these ads[.]").

[130] Email from [redacted], "Re: [Adsense-eng-wat] [Adsense-eng] Re: [Ads-engdirs] Doubleclick Ad Exchange 2.0 - Launched!" (Sep. 19, 2009), GOOG-AT-MDL-010836318, at -318 ("The team has done a great job […] to also go beyond in some important areas, *e.g.* […] integration with Adsense and Adwords.").

[131] Email from [redacted] to [redacted], "Fwd: [adsense-doubleclick-pm] Launch of AWBid: Cross-exchange Buying for AdWords Remarketing" (May 9, 2013), GOOG-DOJ-09916352, at -352 ("I am happy to announce the launch of AWBid in a closed beta in North America. AWBid stands for AdWords Bidder and enables cross-exchange buying for AdWords remarketing campaigns."); Email from [redacted] to [redacted], "[Launch 127698] AWBid Cross-exchange Pilot for AdWords Remarketing" (Jan. 26, 2015), GOOG-DOJ-14514871, at -871 ("AWBid (AdWords Bidder) will allow AdWords campaigns to buy inventory on third party exchanges via RTB integration.").

[132] Design Doc, "Poirot for AWBid Design Doc" (Sep. 10, 2018), GOOG-DOJ-AT-02512863, at -863 ("Motivated by project Poirot, project Poirot for AWBid (A.K.A Marple) aims to provide optimal bidding strategy for GDN advertisers to buy on non-second price exchanges.").

[133] As of 2019, AwBid accounted for approximately [redacted]% of Google Display Ads' total remarketing spend. *See* Presentation, "AWBid Overview" (Sep. 5, 2019), GOOG-DOJ-14298902, at -909 ("[redacted]% of GDA remarketing revenue is from Awbid.").

"Google Display Network (GDN) Auction," the "CAT2" auction, or the "pre-auction"[134]) to determine the bid that Google Ads makes and the advertiser to be allocated the impression if that bid wins.[135] In a first step, Google Ads (like DV360) identifies the advertiser campaigns targeting impressions with the characteristics associated with the bid request and applies its budget throttling algorithm to determine a selection of those campaigns for participation in the auction for the impression.[136] Next, Google Ads converts all of its advertisers' bids into the same **eCPM** basis, which is the expected cost per thousand impressions. For example, to convert an advertiser's value from cost-per-click or cost-per-conversion to eCPM, Google Ads multiplies that advertiser's reported value per click or conversion by a predicted click or conversion rate, which it estimates using a proprietary algorithm.[137] I refer to this eCPM bid as the advertiser's **value** for the impression.[138] Next, Google Ads converts each of these values to a "score,"

[134] Design Doc, "Dynamic Revshare for AdWords on AdX" (Jul. 13, 2012), GOOG-DOJ-13605152, at -152 ("Currently, AdWords runs an auction (aka CAT2 auction, pre-auction) to select the highest bidding creative (or bundle of creatives) to compete in the AdX auction.").

[135] "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -068 ("The GDN auction decides which ads we serve to users for each query."); *see* Presentation, "Auction Overview" (Dec. 2019), GOOG-DOJ-13979867, at -872 to -876.

[136] Overview Doc, "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -068 ("Prior to the auction, the CAT Mixer also applies other filters such as Budget filtering, and so on."); Email from [redacted] to [redacted], "[Launch 205617] Cookie-based budget throttling for GDN advertisers" (Jul. 23, 2018), GOOG-DOJ-15116057, at -057 ("Cookie-based throttling means keying the throttling decision on cookie instead of making an independent decision per query (see design doc for more details). [...] We are now following up to launch on GDN as well.").

[137] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶¶ 4-5 ("Sometimes advertisers use automated bidding products, where instead of providing bids directly, they provide goals (*e.g.*, maximize conversions) and constraints (*e.g.*, budget). Automated bidding systems translate these goals and constraints to an impression value […] To determine which bid Google Ads will submit to AdX, Google computes the impression value for eligible ad candidates so that they can be compared in the same cost units."); "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -070 ("The GDN auction uses the CPM cost type in micros as the standard unit. […] For other cost types, the auction multiplies the bid by a prediction of user interaction.").

[138] Google sometimes also internally refers to the eCPM as the advertiser's value. *See* "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -073 ("Each auction ranks the participants according to a metric named Advertiser Value. Advertiser Value represents the maximum return that an advertiser expects to earn by winning a particular place in the auction."), -074 ("Within GDN, we use MaxEcpm as a proxy for advertiser value").

deducting the Google Ads revenue share to obtain a "net" bid and adjusting in some cases by factors representing the ad's quality and relevance to the user.[139] Finally, Google Ads bids for the impression on behalf of the highest-scoring advertiser (hereinafter, for simplicity, the "highest bidder" on Google Ads).[140]

91. By calculating bids on behalf of advertisers, Google Ads makes it easier for advertisers to participate effectively in auctions for impressions. Even though Google Ads bids into auctions on a per-impression basis, it allows advertisers that mainly care about user engagement—clicks or conversions—to specify a budget and/or values per-click or per-conversion, and Google Ads uses this information to compute bids in the auction on their behalf. With threshold pricing, this hybrid process is bidder-truthful: advertisers are incentivized to report their values per click or per conversion truthfully to Google Ads. In addition, as long as Google predicts engagement rates accurately, this process provides a valuable service, placing bids on behalf of advertisers that are just the same as if advertisers could compute value per impression themselves and pay for each impression according to the threshold pricing rule. For this reason, in the remainder of this report, when I talk about a Google Ads advertiser's payment for an impression, I am referring to the price it would pay if it paid on a per-impression basis (which would be the same, on average, as their actual payments to Google Ads).

92. The process used to determine the Google Ads bid into AdX on behalf of the highest-scoring advertiser has been updated and improved multiple times to increase

[139] "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -071 ("The scoring operation produces the final bid for each candidate by including quality-based adjustments. […] During scoring, the auction calculates bid adjustments and fees, then applies the adjustments to the candidates.").

[140] "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -080 ("The bid that we send to AdX is the actual maximum that our advertiser is willing to pay.").

advertiser profits in response to changing auction dynamics. Initially, Google Ads had a **two-bid policy**, in which it submitted two bids equal to its estimates of the two highest eCPMs that it had estimated among its advertisers (net of a fixed Google Ads revenue share).[141, 142] After adjusting for its revenue share, the two-bid policy replicated the outcome of the second-price auction in which each of its advertisers submitted and determined bids individually. In 2013, in response to other bidders "increasingly thinning their margins to win on volume," Google Ads introduced bid optimization programs, **Dynamic Revenue Share for AdWords** (hereinafter **buy-side DRS**) and its successor, **Project Bernanke**.[143] These programs varied the revenue share charged on a per-impression basis, allowing Google Ads advertisers to win more impressions and allowing publishers to sell more impressions.

93. In 2016, Google Ads launched **Project Bell** which modified its bidding program to protect its advertisers from the publisher tactic called multi-calling, which would otherwise reduce advertisers' surplus. In 2018, Google Ads introduced **Project Elmo** which is a similar program to ensure consistent determination of bids for impressions in

[141] In practice, Google Ads may have submitted one bid with a "minimum payment" field, which specified the least price it would pay if its bid won the auction. All bidders in the AdX auction had the option to use the same field. In practice, the effect of such a field is exactly the same as the effect of submitting two bids, so I do not distinguish them in the remainder of this report. *See* Design Doc, "Call-out-Proxy: 1-bid, 2-bid, and min-bid" (Feb. 24, 2012), GOOG-DOJ-03366145, at -145 ("Engineering suggestion: Each ad-network transfers a single bid, with an added field: a minimum price that needs to be paid if its bid wins."); Launch Doc, "RPO Exemption Policy V2 Launch Doc" (Nov. 14, 2017), GOOG-DOJ-13212948, at -948 ("The current policy (b/18573816) exempts a buyer network on a specific ad query if we see either more than one open auction bid submitted by the network, or a single bid that specifies a minimum payment.").

[142] I use the term "two-bid policy" to describe the pre-2013 Google Ads bidding program. I note that Google Ads continued to submit two bids into AdX auctions during its later bid optimization programs, but these bids were calculated differently (and, in Project Bernanke, the low bid was sometimes zero).

[143] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227.

the face of publisher multi-calling and bid duplication tactics from non-Google exchanges.[144]

94. In 2019, Google Ads further revised its bidding program to make it easier for advertisers to bid in response to AdX's transition to a Unified First Price Auction, with this most recent bidding program called **Alchemist**.

95. I discuss these Google Ads bidding programs and their benefits for advertisers in detail in Section IV, where I also respond to allegations made by Plaintiffs about these programs.

3. Google Ad Manager

a) How Publishers Use GAM

96. **Google Ad Manager (GAM)** is a supply-side platform that includes publisher ad server functionality and the real-time auction capabilities of an ad exchange.[145] The ad server capabilities were previously called **DoubleClick for Publishers (DFP)**, and the ad exchange capabilities were previously called **Ad Exchange (AdX)**. These features were integrated into GAM between 2014 and 2018.[146] In this report, where historically appropriate or where the distinction between the ad server capabilities and real-time

[144] Email from [redacted] to [redacted] "[Launch 205617] Cookie-based budget throttling for GDN advertisers" (Dec. 6, 2018), GOOG-AT-MDL-015521456, at -456 ("Launch Date [...] 2018-11-19").

[145] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ft. 1.

[146] Sridhar Ramaswamy, "Introducing simpler brands and solutions for advertisers and publishers," Google Blog (Jun. 27, 2018), https://blog.google/technology/ads/new-advertising-brands/ ("[W]e've been working to bring together DoubleClick for Publishers and DoubleClick Ad Exchange in a complete and unified programmatic platform under a new name—Google Ad Manager."). Originally, this merged entity was internally called DRX (DoubleClick Reservations and Exchange). *See* Product Requirements Doc, "The 2015 Unified/Integrated DRX Query Tool (a.k.a Project Brundlefly)" (Feb. 5, 2015), GOOG-TEX-00048091, at -093 ("In July 2014, the DoubleClick for Publishers (DFP) and DoubleClick Ad Exchange ads products were merged to form a new product called DoubleClick Reservations & Exchange (DRX). The underlying goal of DRX is to unify the various features offered by both DFP and AdX together to create a more streamlined product experience.").

auction capabilities of GAM is particularly important, I use the older names, DFP and AdX.

97. Publishers include code snippets (called **Google Publisher Tags** or **GPTs**) in their web pages to trigger a **call** or **ad request** to GAM.[147] This ad request contains information about the impression opportunity, including the URL of the website, any available user-related information, and characteristics of the impression opportunity (including its size and location on the page).[148] After an ad request is received, GAM uses decisioning logic to determine which ad to serve.[149] Publishers configure this decisioning logic in the GAM interface using various controls, including **line items**.[150]

98. Line items are GAM's way of encoding information about sources of advertising demand.[151] **Guaranteed line items** are the line items used for campaigns for which the advertiser and publisher have a directly negotiated contract.[152] Guaranteed line items contain information about the advertiser's campaign goal (*e.g.*, number of impressions,

[147] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -487 to -488 ("GPT is a Javascript library that publishers use to tag their web pages so they can talk to Google Ad Manager backend.").

[148] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -488 ("The Ad Request contains information about the impression: URL of the site[,] Browser User Agent[,] Slot parameters (Ad Unit, size, key/value pairs)[,] etc. Ad Request also contains user-related information like Cookies, User IDs, etc, that can be later at the backend matched to user demographics and behavior profiles, audience segments, etc.").

[149] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -492 ("When AdManager receives an ad request, it goes through an ad selection process to pick the Line Item or Yield Group that will serve.").

[150] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -490.

[151] *See* Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -490 ("Line Items represent campaigns / campaign elements within AdManager.").

[152] Google, "Line item types and priorities," Google Ad Manager Help (accessed Jan. 7, 2024), https://support.google.com/admanager/answer/177279?hl=en ("Guaranteed line items […] Standard […] Use this line item type for directly sold campaigns when your buyer wants a specific number of impression[s] to serve.").

clicks), campaign duration, campaign priority, frequency, and targeting criteria.[153] **Non-guaranteed line items** can be used to represent sources of demand for **remnant impressions**, which are impressions not allocated to guaranteed contracts. Those sources of demand include ad networks and non-Google ad exchanges.[154] Until the introduction of EDA in 2014 (see Section III.D.3.e below), guaranteed line items were always prioritized over non-guaranteed line items.[155] Other types of line items can be used to trigger calls to auctions on AdX.[156] Different line items may be relevant for different ad opportunities, depending on publisher decisions in the GAM interface.

b) Early Online Display Advertising and the Waterfall

99. In the early 2000s, online display advertising impressions were primarily sold via either directly negotiated guaranteed contracts between publishers and advertisers, or non-guaranteed contracts between publishers and **ad networks**.[157] Ad networks typically offered to purchase remnant impressions at a fixed or pre-negotiated price, with no

[153] *See* Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -490 ("Line Items define things line [sic]: Campaign goal (how many impressions, clicks, etc)[,] Campaign duration[,] Campaign priority[,] Frequency capping (how often should it appear for one user)[,] Targeting (where and to whom should the campaign serve)").

[154] Google, "Line item types and priorities," Google Ad Manager Help (accessed Jan. 7, 2024), https://support.google.com/admanager/answer/177279?hl=en ("Any third-party ad network or exchange that provides an appropriate ad tag can be represented by a non-guaranteed line item").

[155] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 12 ("From the launch of Dynamic Allocation in around 2007 until Google introduced Enhanced Dynamic Allocation in 2014, line items determined to be 'guaranteed' on a request were always served without considering AdX, AdSense, or any line items determined to be 'non-guaranteed' on that request.").

[156] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -491 ("There are 3 main Line Item types in AdManager with subtypes that differ by campaign goal").

[157] *See* White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413.

obligation on the publisher to fill a minimum number of impressions.[158] As compared to direct deals, ad networks typically did not give a publisher as much control over the ads that would be placed on its website, or the prices it received for any individual impression.[159] While direct contracts and ad networks continue to play an important role in online display advertising, over time, more sophisticated technologies emerged to allow publishers and advertisers to identify more valuable matches.

100. One technology that emerged early in the online display advertising industry was the capability of ad networks to "passback" an unwanted impression for which they did not have a relevant advertisement, allowing the publisher to offer the impression to other demand sources.[160] This passback capability led to the **waterfall**, in which a publisher specified (using code snippets called "passback tags") a list of demand sources to be contacted sequentially.[161] In a waterfall, the first demand source was offered the

[158] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Publishers usually sell their remaining ad space on a non-guaranteed (or 'pre-emptible') basis through their direct sales channel, as well through their indirect sales channel, which may comprise a handful of ad network partners. […] Non-guaranteed ads […] typically sell at a lower price because of the potential that another buyer will pay a higher price after the initial sale, before the impression is actually delivered; hence the 'non-guaranteed' status. With indirect sales, the CPM is usually fixed, but the number of impressions delivered is not.").

[159] *See* David Kaplan, "On Ad Networks: Pork Bellies, Diamonds, Or The New Direct Marketing?," Forbes (Apr. 8, 2008), https://www.forbes.com/2008/04/08/online-ad-networks-tech-cx_pco_0408paidcontent.html?sh=7414ef02cb8e ("All ad networks are not created equal: If all sides can agree on one thing, it's the need for greater clarity to what's being sold and where it's being placed. […] 'In a lot of cases [in terms of ad nets' handling of remnant, or unsold ad inventory], the buyer doesn't really know what they're getting. And the seller doesn't have any control over price.'"). *See also* AffiliateSeeking.com, "Ad Networks" (captured on Jan. 16, 2008), https://web.archive.org/web/20080116101025/https://www.affiliateseeking.com/list/23000001/1.html (listing some examples of ad network pricing options).

[160] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -503 ("Passback means that when one ad system cannot fill an ad request it passes it back to a different ad system.").

[161] *See* Maciej Zawadziński and Mike Sweeney, "What is Waterfalling and How Does it Work?," Clearcode Blog (Sep. 1, 2016), https://clearcode.cc/blog/what-is-waterfalling/("Waterfalling, also known as a daisy chain or waterfall tags, is a process used by a publisher to sell all remnant inventory. […] Waterfalling gets its name from the waterfall-like process for selling inventory—*i.e.* the demand sources are initiated one at a time, one after another."); Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -503.

opportunity to fill an impression. If that demand source declined to fill the impression (or did not have an eligible ad to serve for the impression), the request was passed back to the next demand source on the publisher's list, with the process repeating until the impression was sold or the publisher's list was exhausted, leaving the impression unsold. The waterfall was a highly configurable process, with publishers free to set the order of consideration and prices for each demand source however they wished.

c) AdX Uses Auctions to Match Publishers and Advertisers

101. **AdX** (formerly known as DoubleClick Advertising Exchange or DoubleClick AdX) launched in 2007 and was designed to "bring Web publishers and advertising buyers together on a Web site where they can participate in auctions for ad space."[162]

102. When AdX receives a call from DFP, it runs an auction on behalf of a publisher.[163] This call contains information about the impression and a floor price for the auction. Prior to 2009, buyers did not vary bids based on real-time information about impressions: instead, if an ad opportunity with certain criteria matching the buyer's stated criteria became available, AdX would run an auction using each buyer's *pre-determined* bid for that type of impression.[164] Since the launch of AdX 2.0 in September 2009, AdX has supported **real-time bidding**.[165] Under real-time bidding, AdX sends bid requests containing

[162] AdX was originally launched by DoubleClick, prior to its acquisition by Google. *See* Louise Story, "DoubleClick to Set Up an Exchange for Buying and Selling Digital Ads," New York Times (Apr. 4, 2007), https://www.nytimes.com/2007/04/04/business/media/04adco.html.

[163] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -513.

[164] *See* DoubleClick, "Ad Selection Specifications for Ad Server Version 14.1" (Mar. 27, 2007), GOOG-AT-MDL-007374059, at -136.

[165] Email from [redacted] to [redacted] "Re: [Adsense-eng-wat] [Adsense-eng] Re: [Ads-engdirs] Doubleclick Ad Exchange 2.0 - Launched!" (Sep. 19, 2009), GOOG-AT-MDL-010836318, at -318 ("The team has done a great job

information about the auction to **Authorized Buyers** (previously known as "AdX buyers" and consisting mostly of ad networks and non-Google DSPs), and triggers a request to Google's buy-side products, Google Ads and DV360, to calculate bids.[166] AdX then runs an auction using the bids that it receives, and it then returns the winning ad to DFP, which serves the ad to the publisher's website.[167] Since the launch of Open Bidding, discussed further in Section III.D.3.g and Section XIII, GAM has also incorporated bids on impressions from non-Google exchanges for publishers that enable it.

103. From its launch until 2019, AdX used a second-price auction to allocate impressions.[168, 169] The highest bid, net of the AdX revenue share, would win the auction as long as that net bid was above the floor price. AdX would then charge the winning bidder its threshold price, which is equal to the higher of the second-highest bid in AdX or the auction's floor price.[170] In April 2015, DFP introduced a feature called **Reserve Price Optimization (RPO)** (also known as **Optimized Pricing**) which helped publishers set floor prices in the AdX second-price auction by increasing some floor prices in bid requests to buyers when RPO predicted—based on historical data on the distribution of

[…] to also go beyond in some important areas, *e.g.* […] Real Time Bidding, and of course integration with Adsense and Adwords.").

[166] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -500 to -502 ("The buyer facing side of Google Ad Exchange is called Authorized Buyers […] Ad Exchange sends a Bid Request to DSPs with this information asking them to bid on the impression.").

[167] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -502.

[168] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 5.

[169] During the period in which DRS v1 and v2 was in effect, AdX would charge the winning bidder a price higher than the second-price in a small minority of auctions. See Section XII below.

[170] As noted above, DRS v1 and v2 could change the price paid by advertisers for some impressions.

bids—that the higher floor prices would increase publisher revenues.[171] I discuss RPO in detail in Section XI.

104. In 2019, AdX transitioned to a first-price auction, which I discuss in more detail in Section XIII below.[172]

105. AdX uses a revenue share model to charge publishers for impressions sold on AdX.[173] AdX keeps a proportion of the total revenue from the sales of the publishers' impressions on AdX in a given calendar month, with this proportion referred to as the **AdX revenue share**. The remaining proportion of the revenue, which is passed on to the publisher, is called the **publisher's revenue share**. In combination with the auction's pricing rule (which determines the price paid by bidders in the auction), the revenue share split determines the total payments AdX makes to each publisher. The baseline AdX revenue share is 20%, but different types of transactions can have different revenue shares, and the revenue share could be changed by contract between AdX and the publisher.[174] Until

[171] Email from [redacted] to [redacted], "Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers" (Nov. 12, 2015), GOOG-DOJ-07235914, at -915 ("Between April and October we launched and improved new systems to dynamically set auction reserve prices for AdX sellers."); Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -154 ("Optimized pricing in the Open Auction […] formerly known as Reserve Price Optimization (RPO)").

[172] Sam Cox, "Simplifying programmatic: first price auctions for Google Ad Manager," Google Ad Manager (Mar. 6, 2019), https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/ ("in the coming months we'll start to transition publisher inventory to a unified first price auction for Google Ad Manager."); Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 7.

[173] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -514 ("Revenue share is the pricing model for Ad Exchange.").

[174] *See* Google, "Google Platform Services Terms and Conditions," Google (accessed Sep. 27, 2023), https://www.google.com/google-ad-manager/platform/terms/, at Section 4.1.a; Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -514 ("Baseline revenue share is 80/20 which means that of every dollar an advertiser pays (Gross value), 80 cents go to publisher and 20 cents go to Google. […] Different types of transactions might have a different revenue share and this might be negotiable during contracting phase.").

August 2015, a publisher would receive the same, fixed revenue share on each impression won by AdX.[175] After the introduction of **Dynamic Revenue Share for AdX** (hereinafter **sell-side DRS**) in August 2015, and until AdX transitioned to a Unified First Price Auction in 2019, AdX allowed the revenue share to vary on individual impressions to sell more total impressions and to increase the overall returns to publishers, as I discuss further in Section XII below. AdX does not generally charge fees to advertisers.[176]

106. The revenue share model helps to align the interests of Google and publishers. Each publisher and Google receive a fixed proportion of the publisher's monthly sales revenue on AdX, so that both parties share in any additional revenues when the total volume or the average price per impression of sales on AdX increases. This degree of alignment allows publishers to safely delegate certain decisions to Google and to benefit from Google's efforts to make its marketplace thicker and more efficient.

d) Dynamic Allocation Benefits Publishers and Improves Efficiency

107. In 2007, DoubleClick introduced **Dynamic Allocation (DA)** in DFP.[177] Using DA, publishers could configure DFP to ensure that the impression was sold on AdX only when an AdX buyer was willing to pay *more* than the amount the publisher expected to

[175] "AdX dynamic sell-side rev share (DRS v1) - project description / mini PRD" (Aug. 2014), GOOG-DOJ-03619484, at -484 ("There is a 20% share on all transactions"); Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 28.

[176] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -514 ("Ad Exchange does not impose a buy-side fee").

[177] Dynamic allocation was created by DoubleClick prior to its acquisition by Google. *See* User Guide, "DoubleClick Advertising Exchange User Guide (Beta)" (Mar. 29, 2007), GOOG-DOJ-AT-01133273, at -277 ("Dynamic allocation for sellers. DoubleClick Advertising Exchange automatically determines how to generate the highest return for every impression by dynamically allocating to the highest paying sales channel.").

receive from any other demand source.[178] In contrast to previous ad allocation technologies (including the waterfall), when impressions were sold to AdX buyers, payments to publishers under DA were determined using an auction process, which ensured that the price of each impression was determined by bidders' demands.

108. DA worked as follows. For each remnant line item in DFP (representing, for example, an ad network), publishers would configure a **value CPM** (sometimes called a **static bid**), which determined how that remnant line item would compete with AdX demand under DA.[179] Publishers could set the value CPM for each remnant line item as they pleased and could create multiple line items with different targeting criteria for the same demand source to allow different value CPMs to be used for different categories of impressions.[180] Given these value CPMs, DA used a two-step process to allocate remnant impressions.[181] First, it would identify the eligible[182] non-guaranteed line item with the highest value CPM, called the **DFP booked price** for the auction. In the second step, AdX would run a second-price auction among its demand partners, with a floor price that was at least as

[178] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Instead of randomly rotating other ads into an ad slot, DoubleClick Ad Exchange uses Dynamic Allocation, rotating in higher-paying ads from ad networks and other third-party media buyers when the net CPM they provide to the publisher is higher than what has been booked directly into the ad server.").

[179] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 11; Google, "Value CPM," Google Ad Manager (accessed Oct. 15, 2023), https://support.google.com/admanager/answer/177222?hl=en("The value CPM (cost per thousand impressions) is an amount you specify to help Google Ad Manager estimate the value of campaigns. The amount entered in the 'Value CPM' field serves two purposes: 1. It's used in revenue calculations for impressions served. 2. When a value CPM is defined for remnant line items, the value CPM is used for competition in dynamic allocation and First Look instead of the 'Rate' value.").

[180] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 11.

[181] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 10.

[182] A line item was "eligible" if the impression met the terms of the non-guaranteed contract between the publisher and the demand source (*e.g.*, based on targeting criteria).

high as the DFP booked price.[183] If a buyer on AdX was willing to pay more than the floor price, it was awarded the impression and paid the auction's clearing price (adjusted for the AdX revenue share); otherwise, the impression was allocated to the best non-guaranteed line item (which might then passback the impression to other demand sources, as in the waterfall).[184]

109. DA benefited publishers and improved efficiency, increasing the total value of online display advertising to advertisers. If publishers set a value CPM for each line item that was at least as large as the expected return from allocating an impression to that demand source, DA was a *risk-free improvement in their expected revenue*: DA would assign an impression to AdX only if it could pay more than the return the publisher expected from any other demand source. Advertisers also benefited from the ability to make higher bids for impressions they valued highly and lower bids for impressions they valued less in the AdX auction, allowing them to focus their spending on impressions they valued the most. These benefits of DA were amplified when AdX transitioned to real-time bidding in 2009.

110. I provide a more detailed analysis of the benefits of DA in Section VIII.

[183] If the publisher had otherwise set a floor price for the impression that was higher than the DFP booked price, that floor price would apply.

[184] Presentation, "Understanding the AdX Auction" (Oct. 2014), GOOG-DOJ-12443562, at -582 ("If we decide to call other network and they have nothing, they can pass it back"); Vijay Sivasubramanian, "Help me with Waterfall setup and understanding Passback Tags," Google Ad Manager Help (Mar. 6, 2019), https://support.google.com/admanager/thread/2065360/help-me-with-waterfall-setup-and-understanding-passback-tags?hl=en("You need to give the passback tag which you generated in GAM to the respective network partner. They will place that passback tag as backup ads. So that it will call next ad partner when you target the passback ad unit."); Stack Overflow, "What are advertising passback tags and common implementation" (Feb. 19, 2014), https://stackoverflow.com/questions/19112923/what-are-advertising-passback-tags-and-common-implementation ("When your primary ad-network doesn't have anything to serve (for example, there isn't an advertiser willing to pay a high enough CPM), you can send that impression back to your default advertiser's tags to serve. […] Passback tags allow you to implement a so-called waterfall model").

e) Enhanced Dynamic Allocation Benefits Publishers and Improves Efficiency

111. In March 2014, DFP introduced **Enhanced Dynamic Allocation (EDA)**, which selected the impressions that would be used to serve guaranteed line items in a way that could maximize publisher revenue.[185] Prior to EDA, publishers would first determine whether an impression would be used to fulfill a direct contract, without checking on the bid available from AdX and other remnant demand sources. This procedure left money on the table by allowing an impression to be assigned to a guaranteed contract when it could have otherwise attracted a high price from remnant demand, while another eligible impression not assigned to the guaranteed contract either went unsold or attracted only a low price from remnant demand.

112. Google engineered EDA carefully to avoid such losses. It first estimated the distribution of bids using data from past impressions that were eligible to fulfill the guaranteed contract. It used that distribution to compute a **temporary CPM** for each guaranteed line item, which Google described (correctly) as the opportunity cost of not assigning an impression to the guaranteed contract.[186] The temporary CPM of the best eligible guaranteed line item, called the **EDA price**, was factored into the floor price in the AdX auction. Setting the EDA price as the opportunity cost meant that just enough impressions would be set aside to ensure that guaranteed contracts would be fulfilled.[187] If Google's

[185] Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -161 ("Generally availability roll-out began on 3/3/2014.").

[186] *See* Google, "Delivery basics[:] Ad competition with dynamic allocation," Google Ad Manager Help (accessed Jan. 8, 2024), https://support.google.com/admanager/answer/3721872?hl=en ("The guaranteed line item competes using a temporary CPM or 'opportunity cost' that Ad Manager calculates automatically.").

[187] *See* Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -168 ("Will reservation delivery be affected? Reservation delivery should not be affected. Optimizing revenue while still honoring reservations is the most important goal of this optimization and we've done extensive testing to validate that it works."); Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019),

estimates were exactly correct, remnant demand would buy exactly those impressions for which it was willing to pay the most—all those with remnant bids above the EDA price—and the remaining impressions would be assigned to guaranteed contracts, fulfilling those contracts perfectly. This combination both increases the number of impressions sold and maximizes publisher revenues. EDA also included adaptive procedures to guard against estimation errors. These procedures adjusted the EDA price dynamically over time if too many or too few impressions were being assigned to guaranteed contracts.[188]

113. EDA benefited publishers by increasing the revenues they received from remnant demand without jeopardizing guaranteed contracts. It also ensured that remnant demand advertisers were able to purchase impressions for which their values were the highest, increasing their match values. I discuss the benefits of EDA in detail in Section IX.

f) Emergence of Header Bidding and Integration with DFP

114. Beginning around 2014, publishers began to adopt **header bidding**, a technology that allows publishers to solicit and compare real-time bids from their favored sets of ad exchanges and other demand partners before calling an ad server.[189] To implement header

GOOG-AT-MDL-011687180, at -180 ("Enhanced Dynamic Allocation introduces competition between guaranteed reservations and other demand including AdX and DFP remnant reservations, by allowing AdX or DFP remnant to win over high priority DFP guaranteed reservations if it has a higher price than the opportunity cost (also called EDA price) set by us. The EDA price is calculated in such a way that the DFP guaranteed reservation's delivery goal would not be compromised.").

[188] *See* Google, "DFP and dynamic allocation," DoubleClick for Publishers Help (captured on Sep. 22, 2015), https://web.archive.org/web/20150922150140/https://support.google.com/dfp_premium/answer/3447903 ("The lower a line item's Satisfaction Index (SI) (that is, the more behind schedule it is), the higher the temporary CPM that's passed to Ad Exchange. Therefore, a standard line item that is behind schedule will win often enough to stay on pace to satisfy its goal and pacing settings.").

[189] *See, e.g.*, Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -338 to -339; Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13.

bidding, a publisher inserts a snippet of code in the header of its web page.[190] That code calls participating ad exchanges or other demand partners (either directly or via a server) to submit bids into an auction run on the browser or server, typically in a first-price format.[191] A publisher could directly allocate the impression to the demand source with the highest header bid, but many publishers sought to do better by integrating header bidding with their publisher ad server.[192] Integrating header bidding with DFP allowed publishers to enjoy the benefits of Enhanced Dynamic Allocation and other services on each impression.

115. While DFP was not designed to incorporate real-time bids from non-Google exchanges,[193] publishers incorporated header bidding into DFP by creating separate non-guaranteed line items for each bid value that each header bidding counterparty might provide (*e.g.*, separate line items for "Header Bidding Exchange A" at $1.00, $1.01, $1.02, etc.).[194] I will call those **header bidding line items** for the purpose of this report.

---

[190] *See* Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -338.

[191] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("These Header Bidding auctions are typically first-price auctions."). *See also* Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -943 ("[T]he browser code runs a 1P auction"); Prebid.org, "Prebid.js FAQ" (accessed Dec. 1, 2023), https://docs.prebid.org/dev-docs/faq.html ("Header Bidding is a first-price auction.").

[192] *See* Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -939 ("[redacted]% of LPS [Large Partner Solutions] publishers are using header [bidding] tags.").

[193] *See* Deposition of [redacted] at 71:18-21 (Aug. 12, 2021), GOOG-AT-MDL-007178292, at -363 ("So in the original design of the system, it was not designed to put exchanges in as line items. Line items are designed to represent direct deals or network deals."); Deposition of [redacted] at 50:5-20 (Nov. 6, 2020), GOOG-AT-MDL-007172126, at -176 ("Q: And why was using line items for realtime pricing a risk to Google's AdServer? A: The way the system was built is that line items were always intended to be reservations. There wasn't a concept of using them for realtime pricing. And so we had in mind that publishers would have, you know, possibly thousands of line items and the system was built to scale to that, but with using line items for realtime pricing, which is not what they were designed to be used for, there were ten, sometimes ten times, sometimes 100 times, sometimes 1,000 times more line items than the system was designed to support.").

[194] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 14 ("Up until at least December 2021, the winning bid from the Header Bidding auction was typically used to trigger a specific line item that the publisher had booked within Google's ad server (most commonly a remnant line item) [...] "); Email from [redacted]

After receiving header bids, the publisher's header bidding code would call DFP with the impression and would "trigger" one of the header bidding line items, and, under Dynamic Allocation, the value CPM of that line item could set the floor price in the AdX auction. In that case, if no bid on AdX exceeded that floor price, the impression would be allocated to the header bidder.[195] The value CPMs of header bidding line items—as for all non-guaranteed line items—were determined by publishers, and thus could differ from the winning header bid (and therefore also from the payment the publisher expected to receive for the impression).[196] The huge number of new line items created as a result of publishers adopting header bidding led to a substantial increase in infrastructure costs for Google.[197]

116. In this way, if a publisher configured its website to use header bidding before it offered an impression to AdX, the highest AdX bidder could win an impression if its bid beat the value CPM of the header bidding line item that was triggered by the publisher after

---

to [redacted], "Re: Ultraprio - Increase the ALI for Turner" (Sep. 26, 2018), GOOG-DOJ-11782962, at -962 ("[T]heir HB pricing granularity is in $.01 increments up to $12").

[195] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶¶ 11 ("Under Dynamic Allocation, the Value CPM associated with the best eligible non-guaranteed line item could set the floor price in the AdX auction."), 14 ("[T]he Value CPM of that line item could represent the winning Header Bidding bid as a floor in the AdX auction (prior to September 2019) or as a competing bid in the Unified First Price Auction (from September 2019 onwards).").

[196] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶¶ 14 ("Up until at least December 2021, the winning bid from the Header Bidding auction was typically used to trigger a specific line item that the publisher had booked within Google's ad server (most commonly a remnant line item), and as described above in paragraph 11, the Value CPM of that line item could represent the winning Header Bidding bid as a floor in the AdX auction (prior to September 2019) or as a competing bid in the Unified First Price Auction (from September 2019 onwards)."), 11 ("Some publishers set Value CPMs based on their estimates of what CPM a line item would likely generate (taking into account its historical performance) or based on a fixed price the publisher had negotiated with a particular remnant demand partner. Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons.").

[197] *See* Presentation, "PRD/Strat review: Network health" (Jun. 22, 2017), GOOG-DOJ-06875572, at -589 ("Infra abuse a growing problem: [redacted]X increase in ALI [Active Line Items] since Apr 2016 […] Incremental [redacted]K ALIs (2X limit) increases serving cost by [redacted]% and variable infra cost by [redacted]%").

resolving its header bidding auction.[198] This so-called "**last look**" was not a Google program: it arose as a consequence of the way that some publishers integrated header bidding into DFP using the line item capabilities that DFP (like other publisher ad servers) supported at the time header bidding was introduced.[199] In Section X, I examine the so-called "last look" and show that it was not a source of advantage for Google, contrary to the allegations made by Plaintiffs.

117. Header bidding has benefits and costs for publishers. On the benefits side, header bidding could help publishers increase their online display advertising revenues by collecting more real-time bids for impressions. On the costs side, header bidding has at least four disadvantages. First, it is relatively complex to configure.[200] Second, it often increases page load latency.[201] Third, it makes it more complicated for advertisers to bid optimally because of the possibility of self-competition, where an advertiser drives up the price it pays by unknowingly bidding against itself for the same impression in multiple places.[202]

[198] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 14.

[199] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 17 ("'[L]ast look' was not designed to give AdX an advantage when competing against Header Bidding. It was simply the result of the Header Bidding auction taking place before the AdX auction ran and the way that publishers configured Header Bidding line items to work with Dynamic Allocation.").

[200] One publisher described the time required to integrate an SSP with a leading header bidding wrapper as "20 hours of work" (versus Open Bidding's "20 minutes"). *See* Sarah Sluis, "Google Ad Manager Builds A Bridge To Prebid—But Don't Call It A Two-Way Street," AdExchanger (Apr. 27, 2022), https://www.adexchanger.com/platforms/google-ad-manager-builds-a-bridge-to-prebid-but-dont-call-it-a-two-way-street/. Another industry source compares header bidding to previous ad configurations: "It's not just a little more work, it's probably 100X as much work to traffic for most publishers," *See* Ad Ops Insider, "Header Bidding Explained Step-by-Step" (Jun. 8, 2015), https://www.adopsinsider.com/header-bidding/header-bidding-step-by-step/.

[201] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24. *See also* Vishveshwar Jatain, "Understanding Header Bidding And How To Leverage It," Forbes (Sep. 17, 2019), https://www.forbes.com/sites/forbescommunicationscouncil/2019/09/17/understanding-header-bidding-and-how-to-leverage-it/?sh=332097315c18("Client-side header bidding […] increased page latency because executing auctions takes bandwidth and computing resources."); Pachilakis, M., Papadopoulos, P., Markatos, E. P., & Kourtellis, N. (2019). No more chasing waterfalls: a measurement study of the header bidding ad-ecosystem. In *Proceedings of the Internet Measurement Conference* (pp. 280-293).

[202] *See* Presentation, "Optimal AdX in DFP setup: Best practices, and how to traffic RTA/RTP (header bidding) line items" (Sep. 24, 2015), GOOG-TEX-00000001, at -004 ("[H]eader bidding can make buyers bid against themselves

Fourth, there were reports of payment discrepancies between the bids of header bidding exchanges and eventual payments.[203] I discuss these benefits and costs of header bidding further in Sections X and XIII.

g) Open Bidding Allowed Non-Google Exchanges to Compete in a Unified Auction, Benefiting Publishers and Increasing Platform Thickness

118. Google began testing **Open Bidding** (formerly known as Exchange Bidding, demand syndication, and EBDA) in 2016 and officially launched it in April 2018. Open Bidding allowed publishers to run an auction using real-time bids from multiple ad exchanges, including AdX and competing ad exchanges, within GAM. Some employees described Open Bidding as Google's "answer to header bidding."[204] Open Bidding allows non-Google exchanges to compete for an impression in an "auction of auctions" in which bids from Authorized Buyers, DV360, Google Ads, and other remnant line items booked by publishers in GAM (which might include header bidding line items) compete head to head.

running 2 auctions for every impression."); [redacted]

[203] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24 ("Header bidding is also not transparent because, although the publisher 'accepts' the impression at a certain price, the header bidder may not actually pay the sum indicated in its bid."); Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -955 ("A comparison of HB reports vs DFP reporting showed significant discrepancies [in revenue]"). *See also* James Curran, "Opinion[:] For Publishers, Header Bidding Discrepancies Can Outweigh Revenue Lift," AdExchanger (Jul. 8, 2016), https://www.adexchanger.com/the-sell-sider/publishers-header-bidding-discrepancies-can-outweigh-revenue-lift/ ("Publishers need to create a more realistic calculation of header bidding revenue by factoring discrepancies into their line-item valuations. Some header bidding solutions can cause up to a 50% discrepancy between the publisher ad server impression reports and the impression reports from the programmatic partner. That means a $2 CPM is really a $1 CPM once you account for the adjustments made by the exchange for viewability, verification and performance tracking.").

[204] Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337.

119. Publishers determine which demand sources are eligible to compete under Open Bidding using the GAM interface. Under Open Bidding, a publisher needs to establish a contractual relationship with an Open Bidding partner in order to allow it to compete for the publisher's impressions.[205] When an impression arrives for which Open Bidding is eligible, GAM sends a request to AdX and participating Open Bidding exchanges. The Open Bidding auction design changed several times during testing, as I discuss in detail in Section XIII.

120. The introduction of Open Bidding benefited publishers and thickened the Google online display advertising platform, by simplifying the process of integrating bids from multiple sources, including competing exchanges. Open Bidding is easier to configure and has lower latency than some header bidding alternatives, although many publishers continue to use header bidding exclusively or in combination with Open Bidding.[206] Among its advantages for publishers are simpler configuration and streamlined payments.[207] The introduction of Open Bidding also benefited some non-Google exchanges by reducing the

[205] Google, "Introduction to Open Bidding," Google Ad Manager Help (accessed Oct. 31, 2023), https://support.google.com/admanager/answer/7128453?hl=en("[B]efore a publisher can connect with an Open Bidding yield partner, the publisher must have an established contractual relationship with that partner.").

[206] For a comparison of Open Bidding and header bidding, see Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -440 to -441. For a comparison of timeouts, see Comms Doc, "RTB Timeouts" (Oct. 2019), GOOG-DOJ-15232606, at -609 ("Google's lower bid timeouts should have a slightly better user experience with lower latency"). For usage with header bidding, see Abhilasha, "The Ultimate Guide to Open Bidding for Publishers," headerbidding.co (Jul. 25, 2023), https://headerbidding.co/open-bidding-ultimate-guide/("[I]t is possible to run Open Bidding alongside Header Bidding."); George Levitte, "Improved header bidding support in Google Ad Manager," Google Ad Manager (Apr. 27, 2022), https://blog.google/products/admanager/improved-header-bidding-support-in-google-ad-manager/("[M]any [publishers use] a mix of header bidding and server-side solutions like Open Bidding").

[207] *See* Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -438 ("Eliminate operational inefficiencies such as line item complexity […] Easy to set up, view/analyze reports and unified payments").

publishers' costs of integrating with them, allowing them to sell additional impressions.[208] For example, OpenX reported that "[e]xisting OpenX publisher partners who enabled [Open Bidding] through the OpenX Exchange experienced an average 48% increase in programmatic revenue from OpenX."[209] In Section XIII, I discuss in detail the benefits of Open Bidding for publishers and competing exchanges.

h) Unified First Price Auction and Unified Pricing Rules

121. As discussed in Section III.C.4 above, most exchanges transitioned to first-price auctions between 2017 and 2019.[210] Heterogeneity in the auction formats used by exchanges during this period complicated the implementation of header bidding and Open Bidding, both of which combine the results of auctions on different exchanges that may use different auction rules to sell the same impression. One unfortunate result of this "auction of auctions" process is that the bidder with the highest bid did not necessarily win the impression.

---

[208] *See* Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -438 ("Easy to set up, view/analyze reports and unified payments […] Allows exchanges to respond to RTB call-outs […] Provides integrated reporting and billing for exchange bidding transactions won by 3rd party exchanges"), -441. *See also* Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -775 (Table, "Win Rate of Total DFP Queries"; "EB Other Exchanges," [redacted]").

[209] OpenX, "Google & OpenX Release Study Showing Publisher Partners Experience 48% Revenue Lift Through Google Exchange Bidding Collaboration" (Feb. 15, 2018), https://www.openx.com/press-releases/google-openx-revenue-lift/.

[210] *See* Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -267 (Figure: "Impression Share (US Ad Inventory)"; [redacted]. These figures would not include AdX's later transition to a first-price auction in 2019. *See also* Email from [redacted], "Re: Offering 1st price to publishers?" (Sep. 13, 2017) GOOG-DOJ-05272070, at -075 ("Other exchanges started the migration to 1st price auction recently arguing th[at] it is the best way to integrate it with HB.").

122. GAM transitioned to the **Unified First Price Auction (UFPA)** in 2019.[211] A Google employee described the reason for the transition to the first-price auction on GAM as follows:

> "Publishers typically use our platforms to work with a variety of demand sources, which used to compete under inconsistent rules, and passing through multiple intermediaries, each of which runs its own auction (some first-price, some second-price, and some apparently running strange hybrid mechanisms), potentially transforms bids, and collects a share of revenue. This led to market confusion and inefficiencies; the move to a unified first-price auction was an attempt to move to a simpler, more transparent and sustainable state, improving outcomes for publishers, advertisers, and other ecosystem participants. Of course, first-price auctions are not incentive-compatible, so this required both Google buyers and external buyers to modify bidding algorithms as well."[212]

Under the UFPA, all bidders—including AdX bidders, header bidders, and non-Google exchanges using Open Bidding—compete in the same first-price auction, with no so-called "last look." This change ensured that, for each impression, the winner in the UFPA would be the bidder with the highest bid. Google's buy-side tools implemented bid optimization programs on behalf of their advertiser customers to make it easier for them to bid optimally in the UFPA; there was no need for an advertiser to change its campaign

[211] Sam Cox, "Simplifying programmatic: first price auctions for Google Ad Manager," Google Ad Manager (Mar. 6, 2019), https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/("[I]n the coming months we'll start to transition publisher inventory to a unified first price auction for Google Ad Manager."); Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 37.

[212] Email from [redacted] to [redacted], "Fwd: Join Auction Brown Bag Series: (TODAY @ 12pm PT!)" (Nov. 15, 2019), GOOG-DOJ-AT-00070433, at -433.

parameters in response to the new auction format, as Google's tools did that automatically.[213] I discuss the transition to the UFPA further in Section XIII below.

123. GAM implemented **Unified Pricing Rules (UPR)** at the same time as introducing the UFPA.[214] Under UPR, publishers could set floor prices that varied by properties of the impression and characteristics of the advertiser, but not by the identity of the exchange or demand source of the bidder, ensuring equal treatment of AdX and non-Google SSPs.[215]

124. As I discuss in Section XIV, UPR benefited advertisers and buying tools bidding into multiple exchanges by making it easier for them to determine their bids on different exchanges. Without UPR, publishers could engage in **price-fishing**, which is a close cousin of multi-calling.[216] Price-fishing is a tactic in which some publishers would call the same bidders on different exchanges using different floor prices in an attempt to

[213] *See* Alchemist, discussed in Section IV.C.1.c below (for Google Ads), and Poirot's update for the UFPA, discussed in Section VII.D.4 (for DV360). *See also* Sam Cox, "Simplifying programmatic: first price auctions for Google Ad Manager," Google Ad Manager (Mar. 6, 2019), https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/ ("[A]dvertisers using Google Ads or Display & Video 360 do not need to take any action.").

[214] Jason Bigler, "An update on first price auctions for Google Ad Manager," Google Ad Manager (May 10, 2019), https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/("In addition to impacting how publishers are using floor price rules, changing to a first price auction in Ad Manager requires a change in how our rules function. […] That's why we released a new feature to all publishers globally, called unified pricing rules.").

[215] *See* Comms Doc, "Ad Manager Unified 1st Price Auction" (Sep. 27, 2019), GOOG-DOJ-09714662, at -665 ("Unified Pricing rules will not support the following functionalities that were present in Open Auction pricing rules: Buyer-specific floors: ability to set different floors for different buyers/bidders for a given inventory targeting […] publishers will still be able to: Set per-advertiser floors in Unified Pricing rules"); Google, "Unified pricing rules," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/9298008 ("Advertiser- and brand-specific pricing can be configured in unified pricing rules. They don't apply to remnant line items. Per-buyer and per-bidder pricing are not available."); Jason Bigler, "An update on first price auctions for Google Ad Manager," Google Ad Manager (May 10, 2019), https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/("To maintain a fair and transparent auction, these rules will be applied to all partners equally, and cannot be set for individual buying platforms.").

[216] The main difference between price-fishing and multi-calling is that price-fishing calls the same bidders on different exchanges (with different floor prices), while multi-calling calls the same bidders multiple times (with same or different floor prices). Both tactics can lead to a loss in surplus for an unsuspecting advertiser.

induce them to make unnecessarily high bids to win an impression. When publishers engage in price-fishing, bidders need to account for the possibility that the floor price quoted to them at one exchange may be higher than the floor for the same impression quoted to them on another exchange. That possibility adds costs (because more bids must be submitted) and makes bidding optimally more complicated (because a bidder's optimal bid on each exchange depends on the impression's floor prices on other exchanges). To protect themselves against the possibility that some publishers would price-fish, bidders might be incentivized to reduce their bids on all publishers' inventory, a response that would harm publishers (including those not engaged in price-fishing) and reduce efficiency. Thus, UPR prevented harmful externalities that could reduce the profits of advertisers and publishers alike.

4. Google Responded to Competition by Improving Its Products

125. Google's online display advertising products evolved over time in response to the emergence of new technologies, changes in the behavior of advertisers and publishers, and evolutions in the broader online display advertising ecosystem. I summarize the key changes in Figure 1 above, with buy-side technologies (Google Ads and DV360) below the timeline and sell-side technologies (DFP and AdX) above. Each of these changes improved efficiency or reduced the costs of participating in Google's online display advertising platform, benefiting advertisers, publishers, or both.

IV. GOOGLE ADS BIDDING PROGRAMS: PROMOTING SIMPLE BIDDING AND INCREASING SURPLUS FOR GOOGLE ADS ADVERTISERS

A. Overview

126. Since it started bidding on AdX in 2009,[217] Google Ads has updated its bidding program for the AdX auction several times, with each update designed to increase its profits and the surplus of its advertisers and to respond to changes in the strategies of publishers and other bidders. The following are the key evolutions:

a. Beginning before 2013, Google Ads submitted two bids for its advertisers in each AdX auction that were equal to the two highest values for the impression among its advertisers, adjusted for Google Ads' revenue share (hereinafter the "**two-bid policy**").[218] This two-bid policy was part of a simple, bidder-truthful auction for Google Ads advertisers.

b. In 2013, Google Ads introduced its **bid optimization** programs, **buy-side DRS** and later **Bernanke** to optimize its bidding on AdX and increase the value of impressions won by Google Ads advertisers.[219]

---

[217] Email from [redacted] to [redacted], "Re: [Adsense-eng-wat] [Adsense-eng] Re: [Ads-engdirs] Doubleclick Ad Exchange 2.0 - Launched!" (Sep. 19, 2009), GOOG-AT-MDL-010836318, at -318 to -319 (noting that "[t]he team has done a great job [with AdX 2.0 launch] […] to also go beyond in some important areas, e.g. […] Real Time Bidding, and of course integration with Adsense and Adwords.").

[218] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("In this case, GDN first holds its own auction and submits the leading 2 bids to the AdX auction. Historically, GDN applied a 14% revshare to these bids and AdX applies a 20% sell-side revshare to all bids in their auction.").

[219] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 16, 2013), GOOG-DOJ-04306227, at -227 ("Today we launched GDN Dynamic Revshare - a means for GDN to optimize the revshare we apply to AdX bids."); Launch Details Spreadsheet, Launch 106307 (Aug. 29, 2023), GOOG-AT-MDL-009644018, at cells C1, C2 ("Launch Date [...] 2013-11-11").

c. A 2015 update, **Global Bernanke**, increased the value of impressions won by Google Ads advertisers further by allowing the revenue share collected by Google Ads to vary across different publishers' inventory.[220]

d. In 2019, as AdX replaced its second-price auction with its Unified First Price Auction, Google Ads updated Bernanke for the revised auction pricing rule and called the new program **Alchemist**.[221]

My analysis of data on Google Ads advertisers' values for impressions shows that each of these bid optimization programs increased advertiser surplus for most Google Ads advertisers.

127. Plaintiffs' allegations about Google Ads' bidding programs are marred by their misunderstanding of the operation of these programs. In particular:

[220] Email from [redacted] to [redacted], "[Caqengleads] [Launch 133445] Global Bernanke (+$[redacted] revenue)" (May 21, 2015), GOOG-DOJ-15637938, at -938 ("Global Bernanke is an extension of project Bernanke in which GDN retains a 15% margin on AdX as a whole, while deviating from 15% on individual publishers.").

[221] Design Doc, "The Alchemist (AKA First Price Bernanke)" (Mar. 2019), GOOG-DOJ-14550102, at -102 ("Switching to First Price Auctions (FPA) makes GDN's first price bid optimization (for the winner of Adwords Mini Auctions) a unique problem in many ways: Because of the buy-side infrastructure (like HDMI) which was designed for bidding in a Second Price Auctions (SPA), we want to keep the mechanism truthful and individually rational from the buy-side's perspective and at the same time bid in sell-side's FPA. Furthermore, for each sell-side's publisher, we want to hit a specific (15%) margin. In this document we propose Alchemist: a simple mechanism (and a robust framework) that satisfies all these constraints while maximizing welfare and profit."); Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 22 ("Google updated the Bernanke algorithms in 2019 to be compatible with the Unified First Price Auction. The updated version of Bernanke was sometimes referred to within Google as 'Alchemist.' The update was designed to maintain incentives for Google Ads advertisers to bid their true values even after Google transitioned to the Unified First Price Auction, while continuing to target a similar aggregate take rate for Google Ads as before the transition to the Unified First Price Auction.").

a. Plaintiffs and their experts allege that Bernanke made Google's ad exchange "essentially a third-price auction."[222] This claim and similar ones are wrong. Neither AdX nor Google Ads operated a third-price auction. Bernanke and its variations were buy-side optimizations for Google Ads that determined Google Ads' bids into AdX, and as bidding rules, they could never alter the AdX auction format.[223]

b. Plaintiffs and their experts allege that Google Ads' bid (and revenue share) optimization programs harmed its advertisers as well as publishers,[224] when, in fact, these programs increased the surplus of Google Ads advertisers *without* increasing Google Ads' revenue share, and Google's initial experiments found that the programs increased publisher revenues as well.

[222] Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 359 ("In Project Bernanke, participants believed they were in a second-price auction, but it was essentially a third-price auction"). *See also* Fourth Amended Complaint ¶ 299 ("[Bernanke] switched Google's AdX exchange from a second-price auction to a third-price auction.").

[223] *See, e.g.*, Deposition of [redacted] at 177:13-18 (Apr. 3, 2024) ("No, I wouldn't describe it that way. It's changing the -- its determining the bids that are sent into the AdX exchange auction, but it's not modifying the auction mechanism itself.").

[224] *See, e.g.*, Fourth Amended Complaint ¶¶ 313 ("Bernanke hurt publishers."), 315 ("Bernanke hurt advertisers, too."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶¶ 355 ("Publishers were harmed [...]"), 356 ("Advertisers were harmed [...]"); Expert Report of J. Gans (Jun. 7, 2024), at Section VIII.B.3 ("Projects Bernanke and Global Bernanke harmed publishers by reducing the effectiveness of monetization of their inventory."), Section VIII.B.4 ("Bernanke harmed advertisers by overcharging them in low-demand auctions."); Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.C.2 ("Projects Bernanke and Global Bernanke can lead to a reduction in ad quality as well as revenue per mille for publishers."), ¶ 259 ("Projects Bernanke and Global Bernanke result in lower ROI for GDN advertisers, because GDN advertisers spend more but without generating positive returns (by purchasing impressions at a price exactly equal to their willingness to pay).").

c. Plaintiffs and their experts also allege that Bernanke harmed non-Google buyers and exchanges,[225] when, in fact, they represented ordinary competition, [redacted][226]

d. Plaintiffs' experts further claim that Bernanke and its variants facilitated "collusion" among Google Ads advertisers,[227] but this is a mischaracterization. When Google Ads represented the highest bidder, the price paid by the winning Google Ads advertiser was always at least equal to the highest competing Google Ads bid, when a collusive outcome would entail a lower price. Moreover, and unlike collusion, Bernanke did not suppress output: it caused additional surplus-enhancing transactions to occur.

e. Plaintiffs' experts also speculate that Bernanke may have created inefficiency by "enabling lower-value advertisers to win impressions instead of higher-value

---

[225] Fourth Amended Complaint ¶¶ 589 ("As a result, advertisers bidding through DV360 or other non-Google buying tools are not competing equally with Google Ads for each impression in which Bernanke operates."), 316 ("Bernanke was exclusionary and successfully foreclosed competition in both the exchange and ad buying tool markets"); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 357 ("Other exchanges were harmed because AdX completed manipulated auctions that would not have been completed otherwise, and therefor [*sic*] would have been available to other auctions."); Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.D.1 ("Under Projects Bernanke and Global Bernanke, GDN increased its revenue at the expense of non-Google ad-buying tools."), Section VIII.E.1 ("Projects Bernanke and Global Bernanke did not benefit GDN advertisers, but decreased win rates for advertisers using non-Google ad buying tools."); Expert Report of J. Gans (Jun. 7, 2024), at Section VIII.B.2 ("Bernanke harmed competition in the market for ad buying tools for small advertisers.").

[226] For example, [redacted]

[227] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 233.

ones,"[228] but that theory is not supported by evidence or detailed analysis. For context, many non-Google buying tools submitted just a single bid into the AdX auction,[229] which can incentivize an advertiser using such a tool to submit a bid *higher* than its true value for the impression. Without a program like Bernanke, Google Ads advertisers would have been disadvantaged relative to non-Google Ads advertisers, creating inefficient allocations.

f. Plaintiffs' experts also allege that Bernanke and its variants bestowed an unfair "informational advantage" that "subverted" competitive forces due to the program's confidentiality.[230] These claims, however, overlook the benefits of keeping bidding strategies confidential in auctions, which protects Google Ads' advertisers from exploitation by publishers and other bidders.

B. Google Ads' Two-Bid Policy for AdX

128. Prior to January 2013, Google Ads submitted a **high bid** and a **low bid** into AdX for each impression equal to the values of the highest-scoring advertiser and the second-highest-scoring advertiser in the Google Ads internal auction, adjusted for a fixed Google Ads revenue share (which was typically 14%, and so I assume this to be Google

[228] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 181. *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.C.2 ("Project Bernanke and Global Bernanke can lead to a reduction in ad quality."). Professor Weinberg's arguments only follow "if GDN advertisers tend to display lower quality ads," *id.* at ¶ 247, but he offers no justification for why this presupposition is true. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 759 ("Bernanke enables lower-quality ads to be transacted and displayed on publishers' properties.").

[229] Presentation, "Understanding the AdX Auction" (Oct. 2014), GOOG-DOJ-12443562, at -567 ("[O]ther bidders are allowed to [second-price themselves], but often do not").

[230] Expert Report of J. Gans (Jun. 7, 2024), at ¶¶ 717 ("The higher GDN win rate from the Bernanke program allowed Google to maintain a critical informational advantage over other market participants and, with that advantage, subvert the process of competition."), 725 ("Global Bernanke [...] meant that some publishers may end up paying more to Google but without knowledge of that [...] Thus, the first step in enabling competitive forces to work was subverted.").

Ads' revenue share in the remainder of this section).[231, 232] If the Google Ads high bid won the AdX auction, the impression would be assigned to the highest-scoring advertiser, and that advertiser would pay its threshold price, equal to the clearing price of the AdX second-price auction, plus the Google Ads revenue share.[233, 234] I call this combination of how bids were submitted and prices set the Google Ads "**two-bid policy**."

129. For example, suppose that Google Ads had two advertisers willing to pay $4.00 and $2.00, respectively, for an impression with a floor price of $1.00. Under the two-bid policy, it would submit to AdX a high bid of $3.44 and a low bid of $1.72, with each bid equal to its advertisers' values net of the 14% Google Ads revenue share. If the $3.44 bid from Google Ads was the highest bid on AdX, the Google Ads advertiser would win the second-price auction, and the auction's clearing price would be the larger of the Google Ads' low bid of $1.72 or the highest bid from another bidder on AdX. For example, if

[231] In practice, AdX allowed a bidder to quote a "minimum clearing price" in its bid, which was the minimum price it would pay if it won the second-price auction. There is no difference in effect between a minimum clearing price and a second submitted bid, so I do not distinguish the two elsewhere in this report. *See* Design Doc, "Call-out-Proxy: 1-bid, 2-bid, and min-bid" (Feb. 24, 2012), GOOG-DOJ-03366145, at -145 ("Engineering suggestion: Each ad-network transfers a single bid, with an added field: a minimum price that needs to be paid if its bid wins."); Launch Doc, "RPO Exemption Policy V2 Launch Doc" (Nov. 14, 2017), GOOG-DOJ-13212948, at -948 ("The current policy (b/18573816) exempts a buyer network on a specific ad query if we see either more than one open auction bid submitted by the network, or a single bid that specifies a minimum payment.").

[232] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("In this case, GDN first holds its own auction and submits the leading 2 bids to the AdX auction. Historically, GDN applied a 14% revshare to these bids and AdX applies a 20% sell-side revshare to all bids in their auction.").

[233] "GDN Auction Overview" (Oct. 11, 2014), GOOG-AT-MDL-001094067, at -076 ("The auction price[:] After the winner is selected, the auction converts the cost to a price for each candidate. For example, to compute the price that a MaxCPC candidate must pay in the event of a click, we divide the total cost by the probability of a click: $CPC_i = Cost_i / (pCTR_i * \text{Position Normalize}_{[config, i]})$"), -082 ("AdX bills per impression. When GDN submits a CPC bid, we pay the publisher immediately based on eCPM but we charge our advertiser only if the user clicks the ad. In the logs, the cost_type for these impressions is CPC_TO_ CPM.").

[234] To formalize this mathematically, suppose that v_1 and v_2 are the largest and second-largest values among Google Ads advertisers, and let b_2 be the larger of the second-highest bid on AdX and the floor price on AdX. Under the two-bid policy, Google Ads would make a high bid of $0.86v_1$ and a low bid of $0.86v_2$. If Google Ads provided the winning bid in the AdX auction, it would be charged $p = \max(0.86v_2, b_2)$ and the winning Google Ads advertiser would be charged $p/0.86 = \max(v_2, b_2/0.86)$.

there was another bidder on AdX with a bid of $2.58, the clearing price of the AdX auction would be $2.58 and the winning Google Ads advertiser would be charged $3.00 (equal to $2.58 divided by 86%, to account for the 14% Google Ads revenue share).[235] On the other hand, if the second-highest bid on AdX was the Google Ads low bid of $1.72, the clearing price of the auction would be $1.72, and the winning Google Ads advertiser would be charged $2.00.

130. This two-bid policy had four important effects. *First*, it ensured that Google Ads received a fixed proportion of the revenue generated by the sale of each impression. *Second*, because it charged the advertiser a threshold price, it ensured that the combination of the Google Ads internal auction and the subsequent AdX auction was a bidder-truthful process. This meant that advertisers only needed to determine their values for impressions (or their other campaign goals), and could do no better than to allow Google Ads to choose bids on their behalf. *Third*, the two-bid policy resulted in a market-clearing price for both the advertiser and the publisher. *Finally*, it replicated the outcomes of direct bidding by Google Ads advertisers, in a sense I now formalize.

131. I define a **direct bidding** policy as follows. Under direct bidding:

a. each Google Ads advertiser independently submits a bid to Google Ads with access to the same technical expertise and the same historical information on impressions as Google Ads;

b. Google Ads deducts its revenue share from that bid; and

[235] In this section, to simplify the description of the two-bid policy, I put aside the fact that many Google Ads advertisers pay only when clicks/conversions occur. As I described in Paragraph 91 above, as long as Google's estimates of engagement rates are correct, the average price paid under the pay-per-outcome system is the same as the average impression price given the per-impression bids calculated by Google Ads.

c. for each advertiser, Google Ads submits the result of this calculation as a bid in the AdX auction, and charges a winning advertiser the clearing price of the AdX auction (plus the Google Ads revenue share).

By definition, under direct bidding, no advertiser's bid affects the bid submitted by any other advertiser. The two-bid policy replicates the outcomes that would result from direct bidding.[236]

C. Google Ads' Bid Optimization Programs: Buy-Side DRS, Bernanke, and Alchemist

132. While a two-bid policy maximizes advertiser surplus among all bidder-truthful bidding policies with the same fixed revenue share on each impression, Google Ads identified that it could further improve outcomes for both itself and its advertisers using a different bidding policy that varied its revenue share on an impression-by-impression basis. I call such a program that selects bids to maximize an objective a **bid optimization program** (also known as a buy-side **dynamic revenue sharing program** because, given the bidder's value, one can choose bids indirectly by varying the buying tool's revenue share). Before I describe the bid optimization programs introduced by Google Ads, I first extend the example from Paragraph 129 to show how a program that fixes its average revenue share across a pool of impressions can improve outcomes for Google Ads' customers, while simultaneously increasing Google Ads profits.

133. To illustrate this possibility, suppose that there are *two* impressions being sold on AdX like the ones from Paragraph 129, which have value $4.00 for Advertiser A on Google

[236] To see this, note that an advertiser bidding some amount under direct bidding has the same probability of winning and the same expected payment as under the two-bid policy, so that its optimal bids are the same under the two policies, leading to the same auction outcomes.

Ads and value $2.00 for Advertiser B on Google Ads. As I described above, under the Google Ads two-bid policy, Google Ads would calculate its bids by applying its fixed 14% revenue share to each impression, leading to it submitting bids of $3.44 and $1.72 on AdX. If the best other bid on AdX is $3.80 for the first impression and $1.00 for the second impression, then, under the two-bid policy, the Google Ads bids would only win the second impression at a clearing price of $1.72, which Google Ads would allocate to Advertiser A, charging $2.00, leading to total Google Ads advertiser surplus of $2.00 and Google Ads revenue of $0.28.[237] If instead Google Ads submitted a high bid of $4 and a low bid of $1.00 for both impressions, the Google Ads high bidder would win both impressions, with Google Ads paying $4.80 in total for the two impressions.[238] Google Ads could then charge $5.58 to Advertiser A for the two impressions, leading to increased advertiser surplus of $2.42 (vs. $2.00 under two-bids) and increased revenues for Google Ads of $0.78 (vs. $0.28 under two-bids), while maintaining an average Google Ads revenue share of 14%.[239]

134. In this simplified example in which publishers and bidders left their floor prices and bids unchanged, the decision by Google Ads to maintain its *average* revenue share across multiple impressions while allowing the share to vary for each individual impression

[237] To see this, note that for the first impression, the non-Google Ads bid of $3.80 is the highest, so the high Google Ads bid does not win that impression, but on the second impression, the Google Ads bid of $3.44 is the highest and its bid of $1.72 is the second-highest (and thus the auction's clearing price), so that Advertiser A wins the second impression and pays its threshold price of $2.00. Advertiser A's surplus is then its value minus the price it pays, here $4.00 - $2.00 = $2.00. Google Ads' revenue is $2.00 - $1.72 = $0.28.

[238] To see this, note that the Google Ads bid of $4 is the highest for both impressions, and in the first case, the second-highest bid is $3.80 (leading to a clearing price of $3.80), and in the second case, the second-highest bid is $1.00 (leading to a clearing price of $1.00). The total is then $3.80 + $1.00 = $4.80.

[239] Applying the 14% revenue share to the $4.80 cost of impressions to Google Ads gives $4.80 / 0.86, which is approximately $5.58. Advertiser A's surplus is then $8.00 - $5.58 = $2.42. Google Ads' revenue is then $5.58 - $4.80 = $0.78.

improved Google Ads' profits and the surplus of its advertisers. In general, for *any* fixed set of bids chosen by advertisers and floor prices chosen by publishers, advertisers can benefit from a program that averages revenue shares over a pool of impressions, without changing the overall average revenue share. While in this example, the increase in impressions won by Google Ads led to an equivalent loss in impressions won by other bidders on AdX, *this need not be true in general*: if the example were changed so that instead of the $3.80 being another advertiser's bid, it were the publisher floor price, then no other AdX bidder is harmed, an unsold impression is avoided, and publisher revenue is increased.

135. The foregoing discussion assumed that Google Ads knew its advertisers' values for the impressions and could choose its bids without influencing the incentives for advertisers to report their values for those impressions to Google Ads. It also did not consider the effects of its bidding policy on the floor prices chosen by publishers. As I have emphasized elsewhere, for accurate analysis of any bidding program or auction design, it is important to account for the incentives of both advertisers and publishers. Because these incentives depend fundamentally on each program's designs, I now discuss the evolution of Google Ads' bid optimization programs, before I analyze the effects of each program on advertisers and publishers.

1. The Evolution of Google Ads' Bid Optimization Programs

a) Buy-Side DRS: Reducing Revenue Shares to Win Additional Impressions

136. In January 2013, Google Ads introduced **Dynamic Revenue Share for AdWords** (**buy-side DRS**).[240, 241] Buy-side DRS was introduced partially in response to changes in bidding strategies of other bidders on AdX: a Google engineer observed at the time of its launch that "savvy exchange buyers (especially re-marketing players) are increasingly thinning their margins to win on volume."[242] Buy-side DRS allowed Google Ads to regain volume by reducing its revenue shares on some impressions to allow its advertisers to win additional impressions.[243] Unlike its later bid optimization programs, Google Ads did not increase its revenue share on other impressions.[244] Although this had the effect of slightly reducing Google Ads' overall revenue share, in experiments, that was more than offset by an increased volume of inventory won, so that the program led to an overall increase in profits for Google Ads.[245]

[240] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("Today we launched GDN Dynamic Revshare").

[241] Note that Google had other programs with names like "dynamic revenue share," including on AdX ("sell-side DRS"). For clarity, I will refer to the Google Ads program as *buy-side* DRS.

[242] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227.

[243] Design Doc, "Dynamic Revshare for AdWords on AdX" (Jul. 13, 2012), GOOG-DOJ-13605152, at -152 ("Objective[:] Allow AdWords to dynamically adjust the revenue share it charges its bidders when competing in the AdX auction on a per-query basis […] The goal is to increase match rate of AdWords retargeting ads, while not undercutting Adwords-on-AdX revenue."); Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -350 ("Launch Results […] queries +4.6% […] Adwords effective rev-share: 13.4% (control 14%)").

[244] Design Doc, "Dynamic Revshare for AdWords on AdX" (Jul. 13, 2012), GOOG-DOJ-13605152, at -153 ("In cases with less competition AdWords can keep its current 14% revshare, and in cases where competition is high AdWords can choose to lower its margin.").

[245] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("As a result the leading bid is higher, and more competitive in the AdX auction - enabling GDN to win [redacted]% more impressions. While you might think that using a 0% revshare would eliminate GDN's profit, this is not the case because GDN is only charged the second price in the AdX auction. In fact,

137. Buy-side DRS worked as follows. Google Ads continued to submit two bids for each impression on AdX, but instead of determining the Google Ads high bid by deducting a fixed Google Ads revenue share from the highest scoring bidder's value, Google Ads deducted a *smaller* revenue share (as small as 0%) on some impressions.[246] This had the effect of increasing the Google Ads high bid on a subset of impressions, allowing its advertisers to win additional inventory. Google Ads continued to calculate advertiser payments based on the clearing price of the AdX auction, adjusted for the standard 14% revenue share, unless the result was higher than the advertiser's value for the impression, in which case the payment was capped at the bidder's value.[247, 248]

b) Project Bernanke: Winning More Impressions with a Fixed Average Revenue Share

138. In November 2013, Google Ads replaced buy-side DRS with a more comprehensive bid optimization program called **Project Bernanke**.[249] One of the motivations for Project Bernanke was to expand output by allowing Google Ads advertisers to purchase

---

Dynamic Revshare increases GDN revenue by [redacted]% and profit by [redacted]% on AdX inventory."); Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -348, -350.

[246] Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("Dynamic Revshare changes this logic by applying a revshare of 0% to the leading bid. (The 2nd bid still has the 14% revshare applied to it.)"); Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -347.

[247] To formalize this mathematically, suppose that v_1 and v_2 are the largest and second-largest values among Google Ads advertisers, let b_2 be the larger of the second-highest bid on AdX and the floor price on AdX. Under buy-side DRS, Google Ads would make a high bid of v_1 for some subset of impressions and a low bid of $0.86v_2$. On those impressions, if its high bid won the AdX auction, it would be charged $p = \max(0.86v_2, b_2)$ and would charge a winning Google Ads advertiser $\min(\max(b_2/0.86, v_2), v_1)$.

[248] *See* Design Doc, "Dynamic Revshare for AdWords on AdX" (Jul. 13, 2012), GOOG-DOJ-13605152, at -153; Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -349.

[249] Launch Details Spreadsheet, Launch 106307 (Aug. 29, 2023), GOOG-AT-MDL-009644018, at cells C1, C2 ("Launch Date [...] 2013-11-11").

otherwise unsold impressions, which was around half of all impressions at the time Bernanke was introduced.[250] At a high level, Bernanke achieved this objective by reducing the price that Google Ads paid for some impressions, allowing it to use the savings to bid more on other impressions, thus winning additional impressions for its advertisers and reducing the number of unsold impressions on AdX.[251] Project Bernanke accomplished this by modifying *both* of the bids that Google Ads submitted into the AdX second-price auction: it increased its high bid to win more impressions (as in buy-side DRS), but also decreased its low bid to reduce the AdX clearing price on some impressions.[252] Like buy-side DRS, Bernanke caused the Google Ads revenue share to vary on individual impressions, but Google Ads calibrated Bernanke to ensure that its overall revenue share remained at the 14% target.[253]

139. Project Bernanke used an optimization procedure to determine its bidding strategy into AdX. Under Project Bernanke, Google Ads chose **multipliers** for its two bids into AdX. The **high bid multiplier** led to high bids that were *higher* than those that Google Ads used under the two-bid policy, causing Google Ads advertisers to win additional impressions and reduce the number of unmatched impressions, increasing the publisher's

[250] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("The current match rate on AdX (*i.e.*, queries where there is a winning ad) is about [redacted]%. The primary reason for the low match rate are the reserve prices set by the publisher, which need to be beat for an ad to win the auction.").

[251] Email from [redacted] to [redacted] "[Launch 106307] gTrade: Project Bernanke" (Oct. 18, 2013), GOOG-DOJ-14952787, at -787 ("GDN wins more auctions and generates more revenue at the same average 14% revshare; GDN's advertisers win more auctions and get greater click/conversion volume; and AdX publishers enjoy higher match rate and revenue.").

[252] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("Project Bernanke involves reducing the second price and increasing the first price of the two bids submitted by GDN to the AdX auction in such a way that publishers receive fair payout (e.g. GDN margin remains constant) and GDN profit is maximized.").

[253] Email from [redacted] to [redacted] "[Launch 106307] gTrade: Project Bernanke" (Oct. 18, 2013), GOOG-DOJ-14952787, at -787 ("GDN wins more auctions and generates more revenue at the same average 14% revshare.").

revenue from those impressions and ones for which that bid was second-highest.[254] The **low bid multiplier** led to low bids that were *lower* than those that Google Ads used under the two-bid policy, lowering the AdX clearing price of certain impressions, namely, ones for which Google Ads would otherwise have submitted the two highest bids and both of those exceeded the auction's floor price.[255] The high and low bid adjustments were chosen separately for each publisher to *maximize* the total dollar value of impressions purchased via Google Ads while maintaining an overall Google Ads revenue share target (typically 14%), given the payment rules for Google Ads advertisers.[256]

140. Google Ads performed this optimization procedure weekly using data it obtained from its experiments conducted on a 1% sample of auctions in which it participated in the previous week.[257] It used that data to simulate auctions for different choices of bid multipliers and chose multipliers for each publisher's impressions that maximized the Bernanke objective function, subject to maintaining Google Ads' target average revenue

[254] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("As mentioned above, project Bernanke involves reducing the second price and increasing the first price of the two bids submitted by GDN to the AdX auction in such a way that publishers receive fair payout and GDN profit is maximized. […] We pick various bid multipliers between 1 and 4 and evaluate whether GDN will win that query at each of these bid multipliers. On queries won by GDN at any bid multiplier, the payouts to the exchange and the publisher are then computed for various second bid reductions from 0 to 1."); Presentation, "Project Bernanke: Quantitative Easing on the AdExchange" (Oct. 21, 2013), GOOG-DOJ-12700489, at -493 ("Matched Queries [redacted]

[255] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175 to -176 ("As mentioned above, project Bernanke involves reducing the second price and increasing the first price of the two bids submitted by GDN to the AdX auction in such a way that publishers receive fair payout and GDN profit is maximized.").

[256] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175 ("This is done in such a way that GDN profit is maximized while also ensuring fair GDN payout to the exchange/ publisher. Here, fairness is defined as ensuring the desired margin on the GDN payout. For instance, for non-video requests, this implies retaining only $0.14 on average for every $1 revenue.").

[257] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("In order to gather data for running the auction simulations, a 1% background experiment is run…").

share.[258] Google Ads also introduced an "online safety mechanism" to adjust the multiplier chosen in the case its overall revenue share drifted too far from the target.[259] Initially, the Google Ads revenue share target was applied per publisher.[260] In August 2015, Google Ads launched **Global Bernanke**, which applied the Google Ads revenue share target on average across publishers, while allowing the Google Ads revenue share target to vary to an extent for individual publishers.[261, 262] This additional flexibility in the choice of bid multipliers allowed Google Ads to further increase the total value of impressions won by its advertisers.

141. The optimization procedures described in Paragraph 140 used information of a kind that was also available to other buying tools. Contrary to the characterization by Plaintiffs, Google Ads did not "[rely] on inside information [...] using publishers' unencrypted ad

[258] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("In order to gather data for running the auction simulations, a 1% background experiment is run where every top GDN bid is quadrupled and the second bid dropped. In this experiment, on queries GDN wins, it can be inferred that the second price is the price that GDN needs to beat in order to win the auction. Now, we can determine, if instead of quadrupling the bid, we could have, for instance, tripled the bid and still won this auction.").

[259] Presentation, "Project Bernanke: Quantitative Easing on the AdExchange" (Oct. 21, 2013), GOOG-DOJ-12700489, at -492 ("We implemented safety mechanism to fine-tune bid adjustments when margin drifts away from 14% in each supermixer task. [redacted]

[260] Presentation, "Project Bernanke: Exchange Profit Optimization" (May 20, 2013), GOOG-DOJ-13625417, at -422 ("14% margin across all pubs or per pub? We suggest per pub. Ensures 'fair' payout to each pub").

[261] Launch Details Spreadsheet, Launch 133445 (Aug. 25, 2023), GOOG-AT-MDL-009644112, at cells C2 (noting launch date of 2015-8-12), D2 ("Global Bernanke is an extension of project Bernanke in which GDN retains a 15% margin on AdX as a whole, while deviating from 15% on individual publishers.").

[262] At the same time, Google Ads updated [redacted]. *See* Presentation, "Beyond Bernanke" (Aug. 17, 2015), GOOG-DOJ-28385887, at -894 ("Global Bernanke solves slightly modified optimization problem: [redacted]).

server user IDs."[263] As a Google engineer noted, "[w]e respect GDN-AdX firewall: we only utilize GDN data to optimize bidding strategy. Any AdX buyer can do this."[264] In particular, access to an unencrypted publisher ID was not needed to implement any version of Bernanke: a bidder viewing only encrypted publisher IDs could still track spending by publisher (as in the version of Bernanke before Global Bernanke) using the encrypted ID. Publisher IDs are unnecessary to implement Global Bernanke.

142. In the first version of Bernanke, the prices that Google Ads charged a winning advertiser were determined as follows. On all impressions that it won on AdX at a price lower than the winning advertiser's net bid (after deducting Google Ads' standard 14% revenue share), the winning advertiser paid the larger of the second-highest Google Ads bid and the auction's clearing price (adjusted for the Google Ads revenue share).[265] This is the same price that the advertiser would pay without Bernanke. On the impressions that Google Ads won on AdX at a price that was higher than the winning advertiser's net bid, the winning advertiser paid its bid for the impression.[266] This pricing rule created an

[263] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("It is important to note that in this entire process, we only use information about the GDN bid and the GDN price paid on queries won by GDN. In other words, we do not use any AdX buyer information.").

[264] Presentation, "Project Bernanke: Quantitative Easing on the Ad Exchange gTrade Update" (Oct. 3, 2013), GOOG-DOJ-06842351, at -359.

[265] Presentation, "Project Bernanke: Exchange Profit Optimization" (May 20, 2013), GOOG-DOJ-13625417, at -424 ("Queries GDN was already winning[:] GDN still wins these queries[,] Advertiser cost unchanged, based on second price auction […] Queries GDN wins because of increased bids[:] Advertiser pays first price (same as dynamic revshare)").

[266] To formalize this mathematically, suppose that v_1 and v_2 are the largest and second-largest values among Google Ads advertisers, and let b_2 be the larger of the second-highest bid on AdX and the floor price on AdX. Under Bernanke, Google's high bid was βv_1 (for some $\beta \geq 0.86$) and its low bid was αv_2 (with $\alpha \leq 0.86$). If Google Ads won the AdX auction, it would pay AdX $\max(b_2, \alpha v_2)$, and, under the original Bernanke pricing rule, the winning Google Ads advertiser would be charged $\min(\max(b_2/0.86, v_2), v_1)$.

incentive for advertisers to earn additional surplus on the extra impressions transacted by reducing their bids into Google Ads below their values.

143. In April 2016, Google Ads changed the Bernanke pricing rule for the majority of its advertisers (those using its automated bidding tools) to restore bidder-truthfulness.[267] It did so by charging a winning advertiser its threshold price—that is, an amount equal to the lowest value it could have reported while still winning the impression.[268] This resulted in a bidder-truthful process for Google Ads advertisers, making it simpler for them to configure their campaigns optimally.

c) Alchemist: Updating Bernanke to Optimize Bids for Advertisers in the Unified First Price Auction

144. In September 2019, when AdX completed its transition to the Unified First Price Auction, Bernanke was updated to be compatible with the first-price auction format.[269] The updated program was called **Alchemist**. As I discussed in Section III.C.3.b, in the first-price auction format, it is optimal for a profit-maximizing bidder to shade its bid below its value. Alchemist shades bids into AdX on behalf of advertisers using Google

[267] Launch Doc, "Don't First Price CO on AdX & AWBID" (Apr. 6, 2016), GOOG-DOJ-AT-02467209, at -209 ("Launch has two parts: 1) Do not first-price conversion optimizer ads on AWBID and AdX. Instead charge the minimum price needed to win the query. No change for non-CO ads."); Email from [redacted] to [redacted], "UPCOMING LAUNCH - Please review: [Launch 150218] Removing first pricing in global Bernanke and AWBid DRS + tuning Bernanke thresholds" (Apr. 8, 2016), GOOG-DOJ-15730729, at -729; Presentation, "Auction Overview" (Dec. 2019), GOOG-DOJ-13979867, at -879 ("Preserving incentive compatibility […] Most spend on GDA is from auto bidding").

[268] To formalize this mathematically, suppose that v_1 and v_2 are the largest and second-largest values among Google Ads advertisers, and let b_2 be the larger of the second-highest bid on AdX and the floor price on AdX. Under Bernanke, Google's high bid was βv_1 (for some $\beta \geq 0.86$) and its low bid was αv_2 (with $\alpha \leq 0.86$). If Google Ads won the AdX auction, it would pay AdX $\max(b_2, \alpha v_2)$, and, under Bernanke with the threshold pricing rule, the winning Google Ads advertiser would be charged $\max(b_2/\beta, v_2)$.

[269] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 22 ("Google updated the Bernanke algorithms in 2019 to be compatible with the Unified First Price Auction. The updated version of Bernanke was sometimes referred to within Google as 'Alchemist.'").

Ads. Those bids determine the prices paid by AdX to publishers, but Alchemist charges a winning advertiser its threshold price, which makes the auction process bidder-truthful for advertisers.[270] Like the earlier versions of Bernanke, Alchemist chooses Google Ads' bids in the AdX auction to maximize the total value of impressions won by Google Ads advertisers, subject to maintaining Google Ads' average revenue share and its threshold pricing rule for advertisers.[271] Unlike the original versions of Bernanke, Alchemist optimizes bids for a first-price auction format, rather than a second-price format, determining the optimal bids into the Unified First-Price Auction using experiments similar to those conducted by DV360 in the Poirot program, discussed in detail in Section VII below.[272] This means that under Alchemist, Google Ads shoulders the task of optimizing bids for the first-price auction format—a task that advertisers would otherwise need to attempt to do on their own—while eliminating any need for advertisers to strategize about their reporting to Google Ads.

2. Google Ads' Bid Optimization Programs Benefited Its Advertisers

145. In this section, I establish that each of Google Ads' bid optimization programs increased advertiser surplus for most of Google Ads' advertiser-customers, compared to their

[270] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 22 ("Since Google has transitioned to a Unified First Price Auction, the amount an advertiser pays Google Ads has been determined in generally the same way as before the transition, except that Google Ads began to use minimum-bid-to-win data to determine the amount that the advertiser would need to bid to win the AdX auction (factoring in the Google Ads margin).").

[271] Design Doc, "The Alchemist (AKA First Price Bernanke)" (Mar. 2019), GOOG-DOJ-14550102, at -102 [redacted]

[272] Design Doc, "GDN AdX First-Price Bidding Infrastructure" (Sep. 3, 2019), GOOG-DOJ-15254730, at -735 [redacted]

surplus under direct bidding, as I defined in Paragraph 131 above. In doing so, I fully account for the incentives for publishers to change their floor prices and Google Ads advertisers to change their bids in response to the bid optimization programs. These incentives are overlooked in Plaintiffs' analyses of Project Bernanke.

146. I establish the effects of Google Ads' bid optimization programs on its advertisers in three logical steps. *First*, I consider the effect of each program on advertiser surplus, supposing that the behavior of publishers and Google Ads advertisers remains unchanged by each program. *Second*, I account for the incentive effect of each program on the bids chosen by Google Ads advertisers and show that those effects can only further increase the surplus of Google Ads advertisers. *Third*, I account for the incentives for publishers to increase floor prices, which would tend to offset the increase in advertiser surplus obtained in the first two steps. I use a combination of economic theory and empirical analysis to establish that, for any such program allowing Google Ads advertisers to win more impressions, and even after accounting for publishers' responses, the net effect on Google Ads advertiser surplus (which depends not only on the number of impressions won, but also the prices paid) is positive. Finally, I discuss Google's experimental evidence relating to its bid optimization programs, which support these same conclusions.

a) Fixing Bids and Floor Prices, Each Bid Optimization Program Increased Advertiser Surplus

147. As I have discussed in Section IV.B above, the Google Ads two-bid policy replicated the outcomes of direct bidding and was bidder-truthful, which means that in any single auction, each Google Ads advertiser was incentivized to report its true value for an

impression. I compare outcomes under the various Google Ads bid optimization programs to those that would be obtained under direct bidding, supposing that the behavior of Google Ads advertisers, publishers and other bidders on AdX was unchanged. The following are the effects of each program:

a. *Buy-side DRS, Bernanke,* and *Global Bernanke with its original pricing rule*: These programs all increased the Google Ads high bid into AdX, so that if publishers and other bidders do not change their behavior, advertisers using Google Ads would win all the impressions that they would have won under the two-bid policy and also additional impressions. Under all three programs, for all impressions that Google Ads advertisers would have won in the absence of the bid optimization program, their payments were unchanged, while on the new impressions won as a result of the bid optimization programs, they would pay no more than their values. Overall, this could not reduce the surplus earned by advertisers.[273]

b. *Global Bernanke with a threshold pricing rule* and *Alchemist*: Fixing the behavior of publishers and other bidders, on the impressions that a Google Ads advertiser

[273] Note that the logic for Bernanke must be modified slightly for advertisers using automated bidding strategies involving budgets (because if the bids submitted by Google Ads to AdX were unchanged for those advertisers after Bernanke was introduced, the additional impressions won at a price equal to the bid would reduce the budget, without increasing the advertiser's surplus, and potentially reduce the advertiser's ability to win impressions it would have won in the absence of Bernanke). But Bernanke still benefits such advertisers because—if advertisers do not change their campaign objectives—the automated bidding program could choose a value for each impression leading to the same *net bid* with Bernanke as without it, which would lead to no change in impressions bought or prices (and thus no change in the advertiser's surplus). This means that the *optimal* choices of per-impression value chosen by the automated bidding tool can only lead to an increase in the advertiser's campaign objective, and any change in campaign objectives by a rational advertiser could only *further* increase the advertiser's benefits.

On the other hand, buy-side DRS was not applied to budget-constrained advertisers, so that any additional impressions won as a result of buy-side DRS did not reduce the advertiser's ability to win its existing impressions. *See* Design Doc, "Including Budget Constrained Ads in Bernanke" (Jun. 3, 2015), GOOG-AT-MDL-004555239, at -239 ("Originally, with DRS, we did not include budget constrained ads […]").

would have won in the absence of these programs, the effect of the program is to lower the thresholds that the advertiser must bid to win, and hence to reduce the prices the winning bidder pays, so the surplus it earns on each such impression cannot fall. For any additional impressions the advertiser wins as a consequence of the bid optimization program, the threshold price is never greater than the advertiser's bid, so the additional surplus from those impressions must be positive. This implies that the advertiser earns more surplus with those programs than without them. These programs both have the additional advantage of incentivizing advertisers to bid their value for an impression, making it easier for advertisers to report their campaign objectives to Google Ads, reducing transaction costs for advertisers.

b) If Google Ads Advertisers Re-Optimize Their Bids While Publisher Floors Remain Constant, Advertiser Surplus Can Only Increase Further

148. Under buy-side DRS, Bernanke, and Global Bernanke with the original pricing rule, Google Ads advertisers sometimes paid the amounts they bid on some impressions, which created an incentive for advertisers to lower their bids to increase their surplus on the impressions they won. In contrast, Bernanke with threshold pricing and Alchemist both incentivize advertisers to bid their value for an impression. But whatever an advertiser's *optimal* bids were after the introduction of the Google Ads bid optimization programs, these bids can only result in a *higher* expected surplus than the advertiser obtains under the bids it chooses under direct bidding (which I showed in Section IV.C.2.a was *higher* than the expected surplus it obtains in the absence of each optimization program). This means that—after incorporating any changes in Google Ads

advertisers' bids caused by those programs (but not changes in publishers' floor prices)—each program must have increased expected advertiser surplus.

c) Even Accounting for Likely Reactions by Publishers, Analysis of Google Ads Data Shows that Each Optimization Benefited Most Advertisers

149. Each of the Google Ads bid optimization programs increased some bids submitted by Google Ads for impressions. As a result, a revenue-maximizing publisher would be incentivized to set a higher floor price for each impression.[274] These higher floor prices could, in theory, offset the benefits of the optimization programs for Google Ads advertisers by increasing the prices paid on some impressions and reducing the probability that a Google Ads advertiser wins an impression (compared to the absence of a response from publishers). To assess the combined effects of the bid optimization programs and publishers' responses, I conducted a quantitative evaluation using the Google Ads Log-Level Dataset from January 2024.[275]

150. For my computations, I divided the Google Ads data into **slices**, each containing observations of advertiser values for impressions with the same publisher, inventory unit, and other observable characteristics. I study the [redacted] slices with sufficient data to analyze reliably: those with at least 100,000 bids and $1 of advertiser spend.[276] I provide further details about this dataset and the slices analyzed in the Technical Notes in Section XV.A.2.

[274] This follows since—from the perspective of the publisher—each bid optimization program lowers the overall costs of setting a higher floor price (*i.e.*, it is less likely that there will be no bidder that exceeds the floor price).

[275] Google Ads Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000000000 to -000258388.

[276] These slices were generated using [redacted]

151. I investigate the data by posing the following question: *is it possible that, in some of these slices, publishers could have adjusted their floor prices so that the combined effect of Google's bid optimization and the publisher's adjustment increased Google Ads' win rate without also increasing the surplus enjoyed by Google Ads' advertisers?* If the answer is no, then any Google Ads bid optimization program that increases its win rate in the AdX auction is necessarily also a program that benefits its advertisers.

152. In Theorem 1 in Section XV.A.1, I establish that the answer to the above question depends on the shape of the **distribution** of advertiser values for impressions in each slice. A distribution is a curve showing the proportion of time that an advertiser's value is less than any specified dollar amount. I identify a set of distributions—**the Decreasing Hazard Rate (DHR) distributions**—for which I can offer mathematical proof that if all the slices had DHR distributions, then the answer to the italicized question above is indeed "*no, it is not possible* (see Corollary 1 in Section XV.A.1)." I then plot the distributions of advertiser values from the real-world data for each slice and check how closely a DHR distribution **fits** each of those distributions.[277, 278]

153. The distributions in the Google Ads data fit *very closely* to DHR distributions. In Figure 3, I show the distributions in the real-world data and nearby **fitted DHR distributions** (blue and red, respectively) for two slices. The right panel is the slice (among the [redacted] slices) that is *least well-approximated* by a DHR distribution.[279] The left panel is the slice

[277] A distribution function F with density f has a decreasing hazard rate (DHR) if $f(x)/(1-F(x))$ is a decreasing function.

[278] These fits were calculated using code/gads_bid_optimization_fit.py in my supporting materials. The figures (including Figure 3) are saved in the code/figures directory, with file names prefixed by gads_bid_optimization.

[279] According to a measure of goodness-of-fit discussed in Section XV.A.2.

at the 2nd percentile of goodness-of-fit, so that for 98% of slices, the two curves fit more closely than those in the left panel. For the left panel, the difference between the distribution in the data and the fitted distribution is practically imperceptible: the fitted red curve is almost completely covered by the blue empirical distribution. Even for the right panel—the worst slice among thousands of slices in the data—the fit is still very good, but a small difference is visible between the blue empirical distribution and the red fitted DHR distribution.

Figure 3: Examples of Fitted Distributions At The 2 Percentile (Left) and

154. My results in Corollary 1 imply that, for any DHR distribution of advertiser values, any bid optimization program that increases Google Ads' win rate necessarily also increases advertiser surplus. The corresponding statement when the fit with DHR is close, but not perfect, is that if Google Ads' win rate increases noticeably, for example by 2%, then its advertisers must benefit.

155. In Theorem 1, I make the approximation quantitatively precise so that I can apply it to the actual observed distributions in the Google Ads data. I find that advertiser surplus necessarily increases for at least 95% of the data slices under any bid optimization program that, after accounting for publisher responses, added at least 2% to Google Ads' win rate for that slice.[280] Figure 4 summarizes a set of similar findings, with the horizontal axis corresponding to an increase in Google Ads' win rate and the vertical axis showing the corresponding proportion of slices for which my bound establishes that advertiser surplus must increase.[281] This figure shows that, if a Google Ads bid optimization program led to 2% more matched queries, then advertiser surplus would increase for more than 95% of slices. To put the 2% threshold in context, experiments conducted after the launch of Project Bernanke found that the *total* number of matched queries increased by around [redacted]%.[282]

[280] This result was generated usin [redacted].

[281] Note that my analysis does not imply that other slices of Google Ads advertisers experienced losses due to the bid optimization programs. The DHR condition is a sufficient *but not necessary* condition for increases in the advertiser surplus (given increases in Google Ads' win rates), and the approximations I obtain are not "tight," meaning it is possible that a *larger* proportion of slices also experienced benefits.

[282] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175.

Figure 4: Fraction of Slices in Which Advertiser Surplus Must Increase, as a Function of the Increase in Slice Win Rate

[redacted]

156. The main takeaway from Figure 4 is that, given the value distributions on nearly every slice, that any of Google Ads' bid optimization programs that noticeably increased the Google Ads win rate must have benefited the vast majority of Google Ads advertisers. In other words, for Project Bernanke, Global Bernanke, and Alchemist, Google Ads' incentives were closely aligned with those of its advertiser customers: *any benefits of these programs for Google Ads were accompanied by benefits for its advertisers*.

d) Evidence from Google's Documents and Experiments Support These Findings, and Suggest Benefits for Publishers

157. Google Ads conducted launch experiments related to its bid optimization programs that broadly support my findings and suggest additional benefits to publishers:[283]

a. *Buy-side DRS:* Google Ads estimated that buy-side DRS increased the volume of impressions won by Google Ads by [redacted]%, equivalent to around [redacted] additional impressions won each day, and that most of these would have gone unsold in the absence of buy-side DRS because publishers set floor prices too high for those impressions.[284] These additional transactions that cleared as a result of buy-side DRS increased overall efficiency. While Google Ads' overall revenue share was expected to drop from [redacted]% to [redacted]%, Google estimated at launch that buy-side DRS would increase Google Ads' total revenue by [redacted]%.[285] Later analysis suggested that buy-side DRS increased Google Ads' total spend on AdX by [redacted]% and led to a [redacted]% increase in publisher revenues.[286]

[283] Note that the theoretical effects of the bid optimization programs on publisher revenues are positive or ambiguous. Buy-side DRS led to higher bids from Google Ads (without decreasing any other bids), which could only benefit publishers. The theoretical effect of Bernanke and Global Bernanke on publisher revenue is ambiguous, because it lowered clearing prices for some impressions sold on AdX while increasing the prices of the extramarginal impressions it won. The theoretical effects of Alchemist on publisher revenue are also ambiguous.

[284] *See* Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -348, -350 ("Adwords […] queries: [redacted]%"); Presentation, "Discussion on improving AdX & AdSense backfill" (Apr. 15, 2014), GOOG-DOJ-03876025, at -043 ("Added another [redacted] additional queries per day").

[285] *See* Email from [redacted] to [redacted], "Re: GDN Dynamic Revshare launched today!" (Jan. 17, 2013), GOOG-DOJ-04306227, at -227 ("Dynamic Revshare increases GDN revenue by [redacted]%"); Presentation, "GDN Dynamic Revshare Launch" (Jan. 16, 2013), GOOG-DOJ-02854344, at -348, -350 ("Adwords effective rev-share: [redacted]% (control [redacted]%)").

[286] Presentation, "Project Bernanke: Quantitative Easing on the Ad Exchange gTrade Update" (Oct. 3, 2013), GOOG-DOJ-06842351, at -357.

b. *Bernanke:* Pre-launch experiments found that Bernanke increased the total volume of matched queries on AdX by [redacted]% and increased total publisher revenues by [redacted]%, although these effects would have been felt heterogeneously among publishers.[287] Experiments conducted after the launch found that the total number of matched queries increased by around [redacted]%, as did the total payouts to publishers.[288] Only around [redacted]% of publishers, weighted by revenue, experienced reductions in payouts.[289] The number of clicks and conversions on the impressions won by Google Ads advertisers increased by [redacted]% and [redacted]%, respectively.[290] The updated program, Global Bernanke, further increased publisher payouts (by [redacted]%) and the volume of matched queries (by [redacted]%).[291]

c. *Alchemist:* I have not analyzed (nor am I presently aware of) experiments related to Alchemist, which was introduced at the same time as other large changes to the Google Ads display advertising platform (namely, the Unified First-Price Auction and UPR), making it difficult to identify the effect of Alchemist in isolation.

[287] Presentation, "Project Bernanke: Quantitative Easing on the AdExchange" (Oct. 21, 2013), GOOG-DOJ-12700489, at -493.

[288] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175.

[289] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -180.

[290] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -181.

[291] Email from [redacted] to [redacted], "[Caqengleads] [Launch 133445] Global Bernanke (+[redacted] revenue)" (May 21, 2015), GOOG-DOJ-15637938, at -938 ("Match rate [redacted]% […] Publisher payout +[redacted]%").

D. Responding to Plaintiffs' Allegations about Google Ads' Bid Optimization Programs

1. Bernanke Did Not Alter the AdX Auction Rule and Did Not Deceive Publishers or Advertisers

158. Plaintiffs and their experts allege that Bernanke altered the AdX auction format, claiming that "[w]hen Google Ads submitted the top two bids into AdX, second pricing itself, Project Bernanke worked to turn the second price auction into a third price auction."[292] This claim is a category error, because Bernanke is not an auction rule. As one senior Google engineer described, "[Bernanke] determin[es] the bids that are sent into the AdX exchange auction, but [it does] not modify[] the auction mechanism itself."[293] Google Ads is a *bidder* in the AdX auction and the bidding rules chosen by individual bidders in the auction do not affect the auction format. Neither Google Ads nor AdX *ever* operated a third-price auction. As a consequence, Google did not "deceive[] both publishers and advertisers by converting sealed second price auctions into third price auctions."[294]

159. Bernanke also did not cause Google to "deceive [publishers] into accepting decreased payments."[295] With Bernanke in place, publishers continued to be paid the clearing price of the AdX auction (net of the appropriate revenue share), as was the case prior to the launch of Google's bid optimization programs. Moreover, Google experiments found that its total payments to publishers *increased* overall as a result of Project Bernanke, as was

[292] Fourth Amended Complaint ¶ 552; Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 359 ("In Project Bernanke, participants believed they were in a second-price auction, but it was essentially a third-price auction [...]").

[293] Deposition of [redacted] at 177:13-18 (Apr. 3, 2024) ("No, I wouldn't describe it that way. It's changing the -- it's determining the bids that are sent into the AdX exchange auction, but it's not modifying the auction mechanism itself.").

[294] Fourth Amended Complaint ¶ 560.

[295] Fourth Amended Complaint ¶ 560.

logically necessary for Google to benefit from increased revenues from a program that did not alter its average revenue share.[296]

160. Bernanke also did not "deceive [advertisers] about the price that they were paying to display their ad" and did not "harm[] [advertisers] each time they won an auction that Bernanke or its variations affected."[297] As I discussed in Section III.D.2, an advertiser on Google Ads did not generally pay for each advertisement Google Ads placed on its behalf: instead, Google Ads advertisers typically paid on a price-per-outcome basis. As a result, even before Project Bernanke was introduced, the price paid by a Google Ads advertiser for an "outcome" (*e.g.*, a click on its ad or a sale on its website) would be determined by the clearing price on AdX only *on average* across a pool of impressions. Bernanke worked by averaging not only prices paid to publishers but also *revenue shares* charged to advertisers across pools of impressions, so that—while Bernanke did increase the clearing price paid to publishers for some impressions—on average, advertisers benefited. As my empirical analysis above demonstrates, Bernanke did not "harm" advertisers, and any resulting increase in their win rate improved advertiser surplus for the vast majority of Google Ads advertisers.

[296] Email from [redacted] to [redacted], "[Launch 106307] gTrade: Project Bernanke" (Oct. 18, 2013), GOOG-DOJ-14952787, at -787.

[297] Fourth Amended Complaint ¶ 561.

2. Plaintiffs Omit or Mischaracterize the Benefits of Google Ads' Bid Optimization Programs for Advertisers

161. Plaintiffs and their experts claim that "Bernanke overcharged advertisers,"[298] and that it "harmed advertisers by manipulating and inflating their bids."[299] But Bernanke had two main effects: (1) it decreased the prices paid by Google Ads advertisers on some impressions (namely, those impressions on which Google Ads submitted the two highest bids on AdX), and (2) it used those savings to win additional impressions for Google Ads advertisers, while never charging those advertisers more than their value for the impression. As my empirical analysis in Section IV.C.2.c above shows, the combination of these effects was a benefit to advertisers: any Google Ads' bid optimization program that increased the win rates of Google Ads advertisers also increased advertiser surplus for the vast majority of Google Ads advertisers.

162. Plaintiffs further claim that "Bernanke could route [an advertiser's] bids to less relevant sites and audiences [...] increas[ing] the cost of the [advertiser's] campaign and lower[ing] her return on investment."[300] The first claim is plainly false. Bernanke never altered an advertiser's campaign nor changed what impressions it bid on. The second claim is also incorrect. By design, Bernanke *maximized* the value of impressions won by Google Ads advertisers subject to it.[301] My empirical analysis suggests that the vast

[298] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 772.

[299] Fourth Amended Complaint ¶ 315. *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 12e ("Project Bernanke and Global Bernanke led to an increased win rate for GDN buyers (without improving GDN advertisers' payoffs)."); Expert Report of P. Pathak (Jun. 7, 2024), at ft. 223 ("Project Bernanke does not benefit Google Ads' advertisers").

[300] Fourth Amended Complaint ¶ 315.

[301] *See*"Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175 ("This is done in such a way that GDN profit is maximized while also ensuring fair GDN payout to the exchange/ publisher. Here, fairness is defined as ensuring the desired margin on the GDN payout. For instance, for non-video requests, this implies

majority of Google Ads advertisers must have benefited from Bernanke whenever Bernanke increased the advertisers' win rates.

3. Plaintiffs Ignore the Benefits of Google Ads' Bid Optimization Programs for Publishers

163. Plaintiffs and their experts claim that Bernanke and its variants harmed publishers.[302] But Plaintiffs overlook the fact that Bernanke was a *buy-side program* that benefited Google Ads advertisers. They also overlook experiments conducted after Bernanke's launch that found that total payouts to publishers increased by around [redacted]%,[303] and only around [redacted]% of publishers, weighted by revenue, experienced reductions in payouts.[304] And while some publishers may have lost revenue, Google estimated that the total payout to publishers increased by an additional [redacted]% under Global Bernanke.[305]

retaining only $0.14 on average for every $1 revenue."); Presentation, "Beyond Bernanke" (Aug. 17, 2015), GOOG-DOJ-28385887, at -894 ("Global Bernanke solves slightly modified optimization problem: [redacted]." [Emphasis in Original]).

[302] *See, e.g.*, Fourth Amended Complaint ¶ 313 ("Bernanke hurt publishers."); Expert Report of J. Gans (Jun. 7, 2024), at Section VIII.B.3 ("Projects Bernanke and Global Bernanke harmed publishers by reducing the effectiveness of monetization of their inventory."); Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.C.2 ("Projects Bernanke and Global Bernanke can lead to a reduction in ad quality as well as revenue per mille for publishers.").

[303] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175.

[304] "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -180.

[305] Email from [redacted] to [redacted], "[Caqengleads] [Launch 133445] Global Bernanke [redacted] revenue)" (May 21, 2015), GOOG-DOJ-15637938, at -938 ("Match rate [redacted].

164. Plaintiffs' experts also claim that Global Bernanke reduced publishers' payouts.[306] As support, Professors Gans and Weinberg both cite an internal Google document reporting that some publishers experienced reduced revenue from Global Bernanke.[307] Plaintiffs' experts overlook that the document shows that nearly half of publishers benefited from Global Bernanke. While the other roughly half of publishers saw reduced revenues, those reductions were almost always modest, and even the publishers with the largest revenue reductions still earned more under Global Bernanke than they would have in the absence of any Bernanke program.[308] Combining the effects on both groups, Global Bernanke's total effect on publishers was positive—a [redacted]% increase in payouts.[309] Moreover, Plaintiffs' experts ignore the benefits for Google Ads advertisers described in that same document: the win rate of Google Ads advertisers increased by [redacted]%.[310]

165. Plaintiffs also claim that Bernanke "restricted publisher choice."[311] They argue that "many publishers chose to set higher floors for Google Ads than other demand sources [...] But Bernanke overrode that choice, allowing Google buyers to win at the expense of

[306] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 246 ("In a Project Global Bernanke strategy document, describing an experiment on 10% of the traffic, Google states that [redacted]% of publishers saw decreased revenues under Project Global Bernanke as compared to Project Bernanke, and [redacted]% of publishers saw a decrease between 5 and 10%."); Expert Report of J. Gans (Jun. 7, 2024), at ¶ 767 ("Results of the Global Bernanke experiment show the payout impact breakdown for all publishers. These results show that [redacted]% of publisher payout is negatively impacted by Global Bernanke, with [redacted]% of payout decreasing beyond the 10% margin constraint because of Global Bernanke.").

[307] Launch Doc, "Global Bernanke" (Jul. 26, 2015), GOOG-DOJ-AT-02471194, at -196.

[308] Launch Doc, "Global Bernanke" (Jul. 26, 2015), GOOG-DOJ-AT-02471194, at -194 ("In the experiment, we do see that on occasions publishers do lose slightly more than [redacted]%. [...] However, most of these publishers still earn more than with no [B]ernanke (a future constraint might be to ensure that publishers always make more than they would with no Bernanke in any Bernanke related launch).").

[309] Launch Doc, "Global Bernanke" (Jul. 26, 2015), GOOG-DOJ-AT-02471194, at -196 ("Publisher payout from GDN +[redacted]%.").

[310] Launch Doc, "Global Bernanke" (Jul. 26, 2015), GOOG-DOJ-AT-02471194, at -195 ("GDN win rate [redacted]%[.]").

[311] Fourth Amended Complaint ¶ 314.

non-Google buyers."[312] This is incorrect. In all cases in which publishers set higher floor prices for Google Ads and an increased bid from Bernanke caused a Google Ads advertiser to win the impression, the publisher was paid at least its floor price. If a publisher was unhappy with that outcome, it could always increase its floor price to Google Ads even further. Bernanke generally benefited publishers, as confirmed by Google's experiments.

166. Plaintiffs' experts also allege that Bernanke harmed publishers by increasing the number of "low-quality" ads winning AdX auctions.[313] Professors Gans and Pathak support their claim with anecdotal evidence from the [redacted],[314] citing a Google document in which a Google engineer comments that "[f]rom what I could tell, our auction optimizations (Bernanke) were responsible for [redacted]% of impressions in the [redacted] escalation."[315] It is not clear from that evidence alone how prevalent low-quality ads were and whether changes to Bernanke would affect their prevalence. Moreover, preventing the display of objectionable ads is primarily a filtering challenge rather than a problem associated with any single auction program. While a publisher that was unhappy with the selection of ads winning at one price floor could experiment with increasing its price floors to change the selection of ads, this may not be the most effective approach to addressing ad quality concerns because, with or without a program like Bernanke, it

[312] Fourth Amended Complaint ¶ 314.

[313] *See, e.g.*, Expert Report of J. Gans (Jun. 7, 2024), at ¶ 755 ("Bernanke harms publishers by inflating bids of low-quality ads that would not clear the publisher-set price floor.").

[314] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 760 ("Specifically, they note that Bernanke overrode [redacted] floor prices for [redacted]% of the problematic impressions that led to this issue."); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 166 ("In 2017, Google noted in an internal strategy document that its Bernanke auction optimization '[was] responsible for [redacted]% of impressions in the [redacted] escalation.'").

[315] "Protecting Publishers from Objectionable Ads - Proposal" (May 2017), GOOG-TEX-00782851, at -854.

would still be possible for low-quality ads to win with a sufficiently high bid. As another Google engineer noted, "cpm floors are not the best way to be keeping these ads out."[316] Instead, the cited document identified a range of ways that Google could improve its filtering to avoid low quality ads.[317]

4. Plaintiffs Exaggerate Google Ads' Bid Optimization Programs' Impact on non-Google Buying Tools and Exchanges

167. Plaintiffs' experts repeatedly claim that Bernanke disadvantaged non-Google ad buying tools,[318] but these claims conflate harms to *competition* with effects on *competitors*. Other ad buying tools could and did implement bidding strategies with similar effects as Bernanke, by reducing or omitting the second-highest bid in a second-price auction or by increasing the highest bid.[319] These kinds of bidding programs, which confer significant benefits to advertiser-customers while also increasing a buying tool's profits and win rates, exemplify competition on the merits.

[316] "Protecting Publishers from Objectionable Ads - Proposal" (May 2017), GOOG-TEX-00782851, at -854.

[317] "Protecting Publishers from Objectionable Ads - Proposal" (May 2017), GOOG-TEX-00782851, at -851 ("We propose to build a set of tools to help sensitive publishers minimize their exposure to objectionable ads while avoiding the revenue hit of the 'blunt force' blocks currently used for this purpose.").

[318] *See, e.g.*, Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 180 ("These increases in the win rates for AdX and Google Ads are at the expense of non-Google ad exchanges and ad buying tools."); Expert Report of M. Weinberg (Jun. 7, 2024), at Section VIII.D.1 ("Under Projects Bernanke and Global Bernanke, GDN increased its revenue at the expense of non-Google ad buying tools"); Expert Report of J. Gans (Jun. 7, 2024), at Section VIII.B.2 ("Bernanke harmed competition in the market for ad buying tools for small advertisers.").

[319] *See, e.g.*, [redacted]

168. Plaintiffs and their experts also claim that Bernanke "manipulate[d]"[320] AdX auctions, and, as a consequence, Google Ads "won more AdX transactions and earned more revenue, all at the detriment of rival buying tools."[321] But these claims omit both the observation that other buying tools had incentives to create similar bid optimizations and the evidence [redacted].[322] It also ignores the fact that non-Google buying tools were submitting just a single bid into the AdX second-price auction,[323] and, as a result, advertisers using those tools may have been incentivized to submit bids higher than their values for the impression.[324] As a consequence, if Google did not pursue a program like Bernanke, its advertisers (who had incentives to bid their values) would have been at a disadvantage compared to advertisers using other buy-side tools (who had incentives to bid more than their values). Moreover, Bernanke had a significant expansionary effect on total output by decreasing the number of unsold

[320] Fourth Amended Complaint ¶ 589 ("Project Bernanke, still operates today to manipulate auction results [...] As a result, advertisers bidding through DV360 or other non-Google buying tools are not competing equally with Google Ads for each impression in which Bernanke operates."); Expert Report of J. Gans (Jun. 7, 2024), at ¶ 716 ("[Bernanke was t]he first major project that Google employed to manipulate auction items [...] affecting competition in the exchange market."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 357 ("Other exchanges were harmed because AdX completed manipulated auctions [...]."); Expert Report of J. Andrien (Jun. 7, 2024), at ¶ 54 ("[U]nder Bernanke, GDN manipulates advertisers' bids before sending them to AdX and that Bernanke advantages GDN bidders over non-GDN bidders, counter to Google's representations of equal footing among all auction participants.").

[321] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 750.

[322] *See, e.g.*, [redacted]

[323] Presentation, "Understanding the AdX Auction" (Oct. 2014), GOOG-DOJ-12443562, at -567 ("[O]ther bidders are allowed to [second-price themselves], but often do not").

[324] For example, consider a buying tool that submitted a single bid into the AdX second-price auction on behalf of a group of advertisers, with that bid equal to the highest value reported by the advertisers in that group. If the buying tool charged its advertiser customers the clearing price of the auction, then that policy would lower each advertiser's expected payment for any bid submitted (compared to a counterfactual policy of submitting two bids equal to the values of their two highest-value advertisers) because the second bid submitted by the buying tool can only increase the clearing price of the auction and thus the price charged to the advertiser. Thus, the result of the one-bid policy is a reduction in the expected price paid for each possible bid, which creates an incentive for an advertiser to report higher bids to the buying tool (higher even than its value for the impression) to try to win additional impressions at lower average prices.

impressions—in those cases (as Professor Gans noted[325]), Bernanke does not disadvantage non-Google buying tools.

5. Bernanke was Competitive, Not Collusive

169. Professor Weinberg's claim that Google Ads' Project Bernanke "can be understood as simultaneously facilitating the effects of collusion among GDN advertisers, without their knowledge"[326] is a mischaracterization of Project Bernanke, also repeated by other experts.[327]

170. Professor Weinberg's argument is wrong because Bernanke lacks the characteristic effects of collusion. *First*, when bidding rings participate in a second-price auction, a winning bidder in the ring would pay just enough to beat the bids submitted by other bidders outside of the conspiracy. But under Project Bernanke, a winning advertiser on Google Ads always paid a price at least as high as the second-highest Google Ads bid, which was often higher than the highest bid submitted by other buying tools. *Second*, collusion generally entails reduced output, but Bernanke was a bid-optimization program that increased the number of impressions sold.[328] *Third*, collusion typically harms sellers

[325] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 742 ("Similarly, in the second possible scenario, Bernanke does not impact non-Google advertisers.").

[326] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 233.

[327] *See, e.g.*, Expert Report of J. Andrien (Jun. 7, 2024), at ¶ 43 ("I understand that Project Bernanke and its subsequent iterations, Global Bernanke and First Price Bernanke, facilitated both the effects of collusion among Google Display Network ('GDN') advertisers without their knowledge") (citing Expert Report of M. Weinberg (Jun. 7, 2024), at Section 8). *See also* Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 349 ("In Project Bernanke and a later iteration, Global Bernanke, Google manipulated auctions to increase how often the Google Display Network (GDN) won. For my understanding of this conduct and its basic steps, I partially rely on the expert report of Matthew Weinberg.").

[328] Email from [redacted] to [redacted], "[Launch 106307] gTrade: Project Bernanke" (Oct. 18, 2013), GOOG-DOJ-14952787, at -787 ("GDN wins more auctions and generates more revenue at the same average 14% revshare; GDN's advertisers win more auctions and get greater click/conversion volume; and AdX publishers enjoy higher match rate and revenue.").

in auctions, but Bernanke resulted in *higher revenues* for publishers, as confirmed by Google's experiments.[329]

171. In contrast to Professor Gans' claim that Bernanke "subvert[ed] the process of competition,"[330] Bernanke exemplifies the process of competition among bidding tools and led to benefits for advertisers. Bernanke *maximizes* advertiser returns while maintaining a fixed revenue share, so that *any* tool seeking to maximize advertiser returns while constraining its revenue share in a second–price auction would choose a bidding program similar to Bernanke.[331] Non-Google buying tools adopted similar programs to Bernanke, including omitting or reducing the second bids in the AdX auction[332] and [redacted][333]

[329] *See, e.g.*, "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -175 ("We see good trends in most important metrics. Matched queries, Google revenue, publisher payout increase by about [redacted]%.").

[330] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 717.

[331] For any impression, the price paid to the publisher may be different from the price charged to the advertiser. Adjustments to the two highest bids it submits enables Google to affect both its win rate and its revenue share, increasing the first without distorting the second.

[332] Presentation, "Understanding the AdX Auction" (Oct. 2014), GOOG-DOJ-12443562, at -567 ("[O]ther bidders are allowed to [submit a second bid], but often do not").

[333] [redacted]

6. Confidentiality About Bidding Practices Is Standard, Benefited Google's Customers, and Did Not Impede Competition

172. Plaintiffs' experts allege that Bernanke's confidentiality "allowed Google to maintain a critical informational advantage over other market participants and, with that advantage, subvert the process of competition."[334] However, this claim overlooks the need to keep some bidding strategies confidential in auctions with other strategic participants. Protecting bidding strategies in that way prevents publishers and other bidders from exploiting their bidding information at the expense of lower returns for Google customers. It is also standard in the industry: as one Google employee noted, the introduction of Bernanke "is similar to other third party buyers changing their bidding behaviour which we never announce and is confidential to the buyer."[335] I am not aware of other buy-side tools that make their bidding algorithms publicly available. Indeed, in all my years of studying auctions, I am not aware of any situation where one bidder has been required to disclose its strategies. Doing so would put that bidder at a competitive disadvantage and discourage its participation, reducing auction thickness and lowering returns for sellers.

173. Plaintiffs further claim that "[the] lack of transparency around fees impede[d] other firms from [...] competing with Google by offering the same services."[336] Dr. Chandler goes further and states that Bernanke's "undisclosed Google rule changes [...] made it impossible for auction participants and competing exchanges to understand the rules that

[334] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 717.

[335] "GDN buying change via Bernanke" (Nov. 1, 2013), GOOG-AT-MDL-003995286, at -286.

[336] Fourth Amended Complaint ¶ 351.

governed and applied to auctions run by Google, skewing decision-making and outcomes."[337] These statements are incorrect. Non-Google bidders in the AdX auction do not need all the details of how Google Ads computes its bids for AdX in order to optimize their own bidding strategies. Experimentation—a routine part of the bid optimization process in auctions—can discover optimal bids without access to detailed information about Google's bidding algorithm. Moreover, other buy-side tools implemented programs similar to Bernanke: for example, [redacted].[338] Similarly, [redacted][339]

174. Professor Gans also claims that publishers would be adversely affected by the confidential nature of Bernanke, because "by making pricing non-transparent, publishers would not be alert to any over-payments to Google and would not have the critical first step in being able to initiate a competitive response."[340] But publishers had access to the

[337] Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 360.

[338] [redacted]

[339] [redacted]

[340] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 771.

key data necessary to assess the performance of AdX: their total revenues from auctions on the platform. Moreover, Bernanke, as a Google Ads bidding program, did not change anything about the level of transparency on AdX, a supply-side tool. Publishers are not entitled to bidders' private biddings strategies, and indeed revealing them could be harmful to bidders' outcomes. Google engineers noted this concern with respect to Project Bernanke, worrying that sharing details about how advertiser values determined bids could "be used to harm buyers."[341]

7. Plaintiffs Incorrectly Claim that Google Ads' Bid Optimization Programs Relied on Alleged Advantages Possessed by Google

175. In their Complaint, Plaintiffs allege that Google Ads' bid optimization programs used "inside information to [...] inflate Google Ads' win rate"[342] but buy-side DRS and Bernanke used only data of a kind that was available to other bidders on AdX, namely, the results of experiments conducted on small sets of impressions to optimize bids.[343]

[341] Email from [redacted] to [redacted], "Re: Bid transparency" (Feb. 17, 2017), GOOG-DOJ-13550075, at -076 ("My understanding is that you are asking about whether/when we should agree to share AdWords bids with publishers. [...] With Bernanke boost, our bids do[] not represent a true valuation and can be used to harm buyers."), -075 to -076 ("Even in a 5% sample, publishers can see AdWords high bids due to Bernanke and sur[e]ly it causes more panic on their side that they do not get their fair share. [...] So pub sees that it gets only $[redacted] when AdWords valuation was [redacted] (Note that the real valuation was $1). They cry out louder that AdWords walked out by paying only 20% of their bid (Whereas we paid them [redacted]%). Then they try to become smarter by putting higher reserve price for AdWords ....").

[342] Fourth Amended Complaint ¶ 316.

[343] Presentation, "Project Bernanke: Quantitative Easing on the Ad Exchange gTrade Update" (Oct. 3, 2013), GOOG-DOJ-06842351, at -359 ("We respect GDN-AdX firewall: we only utilize GDN data to optimize bidding strategy. Any AdX buyer can do this.").

There were no barriers to non-Google bidders using experiments to optimize their own bidding strategies,[344] and other bidders did use experiments to optimize their bids.[345]

176. As I demonstrate by example in Paragraph 133, the principle of applying revenue shares on average across impressions (rather than a fixed revenue share on each impression) can benefit advertisers, even in a modest pool of impressions. Google Ads' implementation of Bernanke used experiments conducted on a relatively small sample of daily data (1%), suggesting that other bidders could have implemented similar programs. For example, other bidders could use a larger fraction of daily data, less frequent updates, or more advanced designs that used the existing volumes of data more effectively. Notably, Plaintiffs do not provide any evidence that Google's alleged advantages are so significant that other firms could not create their own bid optimization programs. Moreover, [redacted].[346]

[344] *See* "Bernanke experiment analysis" (Sep. 3, 2013), GOOG-DOJ-13469175, at -176 ("The optimal combination of first bid increase and second bid decrease for each publisher is estimated using AdX auction simulations. In order to gather data for running the auction simulations, a 1% background experiment is run where every top GDN bid is quadrupled and the second bid dropped. […] It is important to note that in this entire process, we only use information about the GDN bid and the GDN price paid on queries won by GDN. In other words, we do not use any AdX buyer information.").

[345] For instance, Verizon's DSP experimented to evaluate its own bid shading program. *See* Zhang, W., Kitts, B., Han, Y., Zhou, Z., Mao, T., He, H., Pan, S., Flores, A., Gultekin, S., & Weissman, T. (2021). Meow: A space-efficient nonparametric bid shading algorithm. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (pp. 3928-36), at 3933. *See also* Choi, H., Mela, C.F., Balseiro, S. R., & Leary, A. (2020). Online display advertising markets: A literature review and future directions. *Information Systems Research*, *31*(2), 556-575, at 562 ("To overcome the econometric issues and to better understand the value of ad spend, both advertisers and researchers are turning attention to randomized field experiments."), 566 ("The evolution of buying and selling practices in display advertising are tightly integrated with the development of intermediaries' enabling technologies (*e.g.*, data collection, data analysis, RTB, real-time ad serving, A/B testing and optimization tools). With such technologies, intermediaries can better match publishers' impressions (consumers) to more relevant advertisers.").

[346] *See, e.g.*, [redacted]

V. PROJECT BELL: GOOGLE ADS' RESPONSE TO THE HARMFUL PRACTICE OF MULTI-CALLING

A. Overview

177. In October 2016, Google Ads launched **Project Bell** to protect its advertisers from a publisher tactic known as **multi-calling**, in which publishers requested bids from AdX multiple times for the same impression.[347] If unaddressed, multi-calling would lower advertisers' profits and harm other platform participants.

178. In their Complaint, Plaintiffs allege that Project Bell affected publishers that do not "give preferential access to AdX,"[348] when, in fact, Bell affected only publishers that called AdX multiple times for the same impression, regardless of whether they partnered with competing exchanges.[349] Plaintiffs' experts do not make substantive allegations about Project Bell.

B. Multi-Calls Harm Advertisers, End Users, and Non-Multi-Calling Publishers

179. **Multi-calling** refers to a tactic in which a publisher (or a supply-side tool on its behalf) makes multiple calls to a demand source for the same impression.[350] In 2016, Google

[347] "Display Ads Quality Launches 2016 - Revenue OKR History" (Apr. 19, 2017), GOOG-AT-MDL-003465605, at tab "2016 Q4 Web & mApps," cells D46, J46 (reporting launch date of "2016-10-26"). *See also* Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -119 ("The objective of this launch is to take steps towards protecting advertisers from price inflation tactics by publishers and exchanges. There are particularly two mechanisms that we try to tackle here (a) multiple call to adwords and (b) unclean auctions on AWBid.").

[348] Fourth Amended Complaint ¶ 311.

[349] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 16 ("Bell v.2 changed Google Ads' bidding behavior only for the publishers that were understood, based on internal experiments, to be calling AdX multiple times for the same potential ad opportunity ('multi-calling publishers').").

[350] A *call* is when a publisher requests that an exchange or DSP bid for an impression. In industry usage, *multi-calling* typically refers to a publisher calling an advertiser multiple times for the same impression. In this section, when I discuss multi-calling, I will focus on the case where a publisher calls AdX multiple times for an impression. Internally, Google also called such publishers "mediating publishers," but this term was *also* used for

found that around 10% of publishers (henceforth **multi-calling publishers**) were multi-calling AdX, leading to an increase in bid requests for Google Ads.[351] Other demand sources also reported a higher bid request volume due to multi-calls and adopted practices to deter those.[352] There are two main reasons that a publisher might choose to multi-call a demand source.

180. *First*, the publisher might hope to use multi-calling with different floor prices to make an unsuspecting advertiser pay more for an impression. This tactic is illustrated in Figure 5 (reproduced from a Google document), and I now describe it using an example of how this tactic can allow a publisher to extract additional revenue from advertisers. Suppose that a publisher knows that an AdX bidder's bid for an impression is $1 for 90% of impressions, and $10 for 10% of impressions (net of AdX's revenue share). However, the publisher does not know which impressions draw the higher bid of $10. If the publisher

publishers who called multiple ad networks or exchanges, and so I avoid this term to avoid confusion. *See* Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -119 to -120 ("In this doc, the term mediation refers to multiple calls.").

[351] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶¶ 12-13 ("Some publishers would call an ad exchange, such as AdX, multiple times for the same potential ad opportunity. For simplicity, I will refer to this practice as 'multi-calling.' […] Some publishers employing multi-calls would set a different floor price for each of the multiple calls made from a single ad exchange for the same potential ad opportunity. For example, a publisher could configure an AdX call with a floor price of $5 for a potential ad opportunity and a second AdX call with a floor price of $4.50 for the same potential ad opportunity."). *See also* Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -120 ("About 10% of the domains on Adx fall in this category.").

[352] *See, e.g.*, Seb Joseph, "DSPs are cracking down on bid duplication," Digiday (May 12, 2020), https://digiday.com/media/dsps-are-cracking-down-on-bid-duplication/ ("Traffic spikes have caused increased costs in processing bid requests, giving already under pressure demand-side platforms extra economic incentive to squash bid duplication. MediaMath is building a new supply chain that doesn't include SSPs that sell duplicated impressions, for example. Last month, The Trade Desk gave all the SSPs it buys impressions from two weeks to stop sending it duplicated requests to take part in the same auction."); Sarah Sluis, "The Trade Desk suppresses bid duplication amid COVID-19 traffic surge," AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge/ ("Publishers that slot the same six exchanges into multiple header-bidding auctions, such as Prebid, Google open bidding and Amazon Transparent Ad Marketplace, send 18 identical bid requests for the same piece of inventory to The Trade Desk. Currently, DSPs often begrudgingly evaluate them all, and find it hard to tell if the impression is exactly the same.").

offers the impression just once using AdX, then its optimal floor price is $1, selling all impressions, but earning no extra revenue on the impressions worth $10. Instead, as illustrated in the internal Google document reproduced as Figure 5, a multi-calling publisher might first call AdX with a floor price of $10, and then, only if the advertiser does not buy the impression on that first call, it could call AdX again with a lower floor price. If the bidder on AdX does not suspect that it will be multi-called and bids truthfully on each call, then it will pay higher prices—sometimes $10—due to these multi-calls. In this example, multi-calling allows the publisher to capture the entire surplus that the advertiser would otherwise earn in a second-price auction.

Figure 5: Multi-calling with decreasing floors.[353]

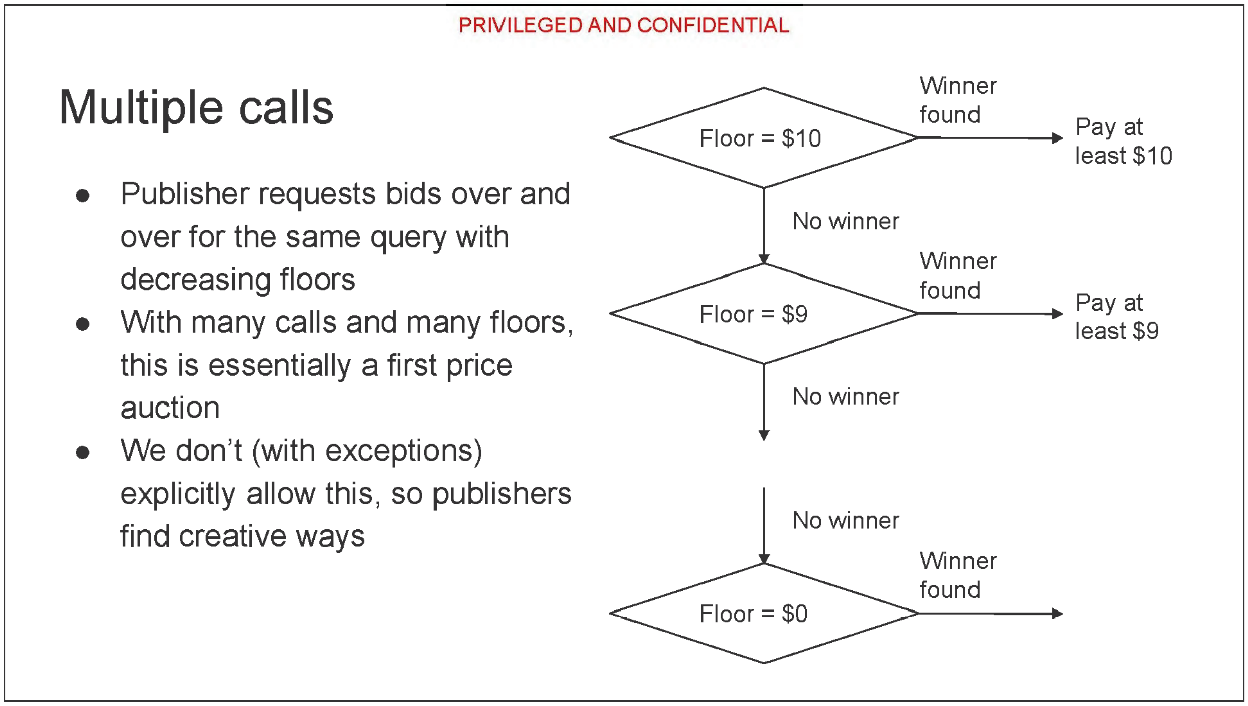

181. Multi-calling with descending floor prices can effectively convert a second-price auction into a **Dutch auction**, in which the seller starts with a high asking price and gradually reduces the price until a buyer is willing to accept (by offering a high enough bid). If a bidder does not adapt its bids to this multi-calling auction format and instead acts as if it is participating in a second-price auction, it ends up paying higher prices for impressions. Instead of bidding truthfully, the profit-maximizing response to multi-calling may require the advertiser to avoid bidding on some calls from the multi-calling publisher (namely, the ones with the highest floor prices). In the example in the previous paragraph, if the AdX advertiser who bids $10 into a second-price auction anticipates multi-calling, it should withhold its bid when the floor price is high and try to win the impression at a lower price when the publisher calls a second time. Deciding when to bid in the presence

[353] Presentation, "Exchange Buying Dynamics" (Aug. 22, 2017), GOOG-DOJ-12848608, at -616.

of multi-calling is complicated by the fact that the advertiser typically knows neither how many times the publisher will call it to bid on an impression, nor the floor prices of those calls, nor how other bidders may be responding to the publisher's multi-calling. For these reasons, determining a bidder's optimal bidding strategy in this setting resembles the problem of bidding optimally in a first-price auction: in both cases, the bidder needs to guess what other bidders are likely to be willing to pay in order to decide when and how much to bid on the publisher's impression. Because it forces bidders to strategize in a way similar to a first-price auction, Google engineers described this multi-calling strategy as "fishing for [a] first price."[354] But multi-calling is more complex and likely to result in worse outcomes than a first-price auction because the bidder must not only make forecasts about other bidders' behavior, but also the behavior of the publisher.

182. *Second,* the publisher may hope that, even with the same floor price, the bidder will change its bid for an impression if called multiple times. Bids for a single impression might change in relatively short intervals of time for a variety of reasons, including the budget pacing or frequency targets that advertisers often employ to govern their bidding.[355] For example, an advertiser may choose not to bid on every impression to avoid depleting its budget in a short period of time or to avoid presenting the same advertisement to the same end user in rapid succession. Multi-calling allows the publisher to undermine these campaign objectives by misrepresenting a single impression as

[354] Email from [redacted] to [redacted], "[Launch 300192] Apps Bernanke for non-compliant publishers" (Jun. 18, 2019), GOOG-DOJ-15119015, at -015. While this quote refers to multi-calling in the context of app inventory, the general principle of "fishing for [a] first price" applies equally to other multi-call settings.

[355] For example, Google Ads offers a frequency capping feature. *See* Google, "Use frequency capping," Google Ads Help (accessed Jul. 24, 2024), https://support.google.com/google-ads/answer/6034106 ("Frequency capping allows you to limit the number of times ads appear to the same person.").

multiple ones. This can allow the publisher to extract higher bids on one of the calls, thwarting the advertiser's campaign goals and reducing efficiency. In response to this tactic, advertisers with frequency caps or budgets are incentivized to bid less or less frequently on impressions in order to achieve their campaign objectives.

183. The net effect is that multi-calling makes bidding more complex and harms efficiency. In the presence of multi-calling, advertisers are unable to tell whether multiple calls from a publisher represent a single advertising opportunity or multiple ones and cannot be sure whether the quoted floor price for an impression is really the lowest bid that the publisher will accept. They need to spend resources to detect and track multi-calling behaviors and adjust their bids strategically, which increases both transaction costs and the likelihood of bidding errors.[356] In the absence of a program that provides these services, the losses caused by multi-calling could lead advertisers to reduce their investments in online display advertising or switch to other buy-side tools that do not bid for impressions offered by multi-calling publishers.[357]

184. Aside from the direct harms caused to advertisers, multi-calling also creates externalities that harm other publishers and users. When advertisers are unsure about which publishers are multi-calling or which calls are duplicates, they are incentivized to reduce their bids into *all* auctions to lower the possible harms of multi-calling. That can reduce efficiency,

[356] *See* Sarah Sluis, "The Trade Desk suppresses bid duplication amid COVID-19 traffic surge," AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge/ ("Publishers […] send 18 identical bid requests for the same piece of inventory to The Trade Desk. Currently, DSPs often begrudgingly evaluate them all, and find it hard to tell if the impression is exactly the same.").

[357] For example, The Trade Desk and MediaMath refused to work with sell-side platforms that made "duplicated requests" for the same impression. Seb Joseph, "DSPs are cracking down on bid duplication," Digiday (May 12, 2020), https://digiday.com/media/dsps-are-cracking-down-on-bid-duplication/.

including by decreasing the overall number of impressions sold, and reduce revenue for publishers, especially those that are not multi-calling. Furthermore, the additional bid requests created by multi-calling inevitably increase costs and latency, which harms all participants. Increased latency harms users, who face slower page load times; advertisers, whose ads may not be seen as a result of the delay; and publishers, whose websites suffer reduced traffic due to poorer user experience.[358, 359]

C. How Bell Identified Multi-Calling Publishers and Protected Advertisers

185. Project Bell identified multi-calling publishers by running experiments on subsets of each publisher's impressions, comparing the number of calls Google Ads received from the publisher if Google Ads bid very high on a subset of impressions (therefore always winning the impression on the first call from a multi-calling publisher) versus the number of calls it received if Google Ads bid very low (therefore always seeing all calls from a multi-calling publisher for an impression).[360] If there is no multi-calling, the number of calls should be the same for both treatments. At the time Bell was introduced, Google

[358] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 15 ("In addition, multi-calling can increase ad latency because it takes time each time the ad exchange is called, an auction is run, and the ad exchange returns an ad.").

[359] Oded Poncz, "Traffic duplication might be a bigger problem than ad fraud," AdExchanger (Jan. 11, 2016), https://www.adexchanger.com/data-driven-thinking/traffic-duplication-might-even-be-a-bigger-problem-than-ad-fraud/ ("Another side effect of bid request duplication is that re-auctioning a bid takes time. In some cases, this could even become apparent to the end user. For example, if an end user is waiting for an interstitial ad to be shown upon opening an application, while behind the scenes there are a few auctions taking place for the same ad space, this will result in the user waiting considerably longer until they're able to access their content. This doubling of ad requests also adds to the overall latency in the ad ecosystem. If all this duplication means an ad takes five seconds to load, that makes for a poor consumer experience.").

[360] Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -120 ("In summary, we run two experiments - one where we bid very high and the other where we bid very low. If the query count on Adx drops by over [redacted]% from the low to high bid experiments over a week period, we consider the domain to be a mediating domain.").

Ads' tests revealed that around [redacted]% of publishers on AdX were using multi-calls.[361] Contrary to Plaintiffs' allegations, Project Bell only affected publishers who made multiple calls *to AdX* for an impression. Making multiple calls to different exchanges before calling AdX did not trigger Bell.[362, 363]

186. Project Bell applied three treatments to protect Google Ads advertisers and discourage multi-calling by publishers.[364] Each treatment increased advertiser surplus by reducing the likelihood that Google Ads would bid above the highest floors chosen by multi-calling publishers.

187. The first treatment imposed a maximum, or **cap**, on bids to multi-calling publishers.[365] This helped protect Google Ads advertisers from paying too much to acquire these impressions. If an advertiser was aware of the publisher's multi-calling behavior, it could cap its own bid on impressions from that publisher, but it would be difficult for an

[361] Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -120 ("About [redacted]% of the domains on Adx fall in this category.").

[362] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 16 ("Bell v.2 changed Google Ads' bidding behavior only for the publishers that were understood, based on internal experiments, to be calling AdX multiple times for the same potential ad opportunity ('multi-calling publishers'). This launch did not directly change Google Ads' bidding behavior for any other publishers' ad opportunities. For example, Bell v.2 did not affect Google Ads bids on the inventory of publishers that called other exchanges before AdX, but did not call AdX multiple times, even if those other exchanges in turn called Google Ads.").

[363] I have reviewed documents (*e.g.*, Presentation, "Mediation detection and GDN bidding" (Mar. 2, 2015), GOOG-DOJ-06563186, at -194 to -201) in which Google considered applying these treatments to publishers who used "mediated single-calls," *i.e.*, publishers that called another exchange before AdX. I understand these plans were never implemented. *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 16.

[364] Presentation, "Mediation and non-second price auction: Detection and treatment" (Jun. 30, 2017), GOOG-DOJ-09875989, at -6008 to -6012.

[365] *See* Email from [redacted] to [redacted], "[Launch 188932] More aggressive treatment for mediation publishers" (Apr. 21, 2017), GOOG-DOJ-15208416, at -416 ("The only difference between this launch and [redacted] is the treatment for these mediation publishers in Adx. We use more aggressive treatment in this launch, [redacted]."). The level of this cap varied over time.

advertiser to detect that behavior and reduce bids accordingly.[366] Google Ads provided a service to its advertisers that would likely be more accurate and less costly than what many advertisers could implement on their own.

188. The second treatment was to disable Bernanke for multi-calling publishers.[367] By increasing Google Ads' high bids, Bernanke increased the variability of Google Ads' bids and thus the potential profitability of multi-calling for publishers.[368] Consequently, Bell disabled Bernanke for these publishers and reverted to submitting its two highest bids with a standard 14% revenue share to avoid incentivizing multi-calling from publishers and to protect Google Ads advertisers and non-multi-calling publishers.[369]

189. The third treatment was to stop buying from multi-calling publishers through AwBid, the Google Ads feature that allows advertisers to purchase some types of inventory through non-Google exchanges.[370] For a given impression, publishers making multiple calls to AdX might also make duplicated calls to other exchanges, which would be more difficult

[366] *See* Seb Joseph, "To reduce auction duplication, buyers start to enforce sellers.json," Digiday (Oct. 16, 2019), https://digiday.com/media/reduce-auction-duplication-buyers-start-enforce-sellers-json/ ("The impasse made it hard to spot duplicate bid requests from the same exchange for the same impressions."). *See also* Heymann, B. (2020). How to bid in unified second-price auctions when requests are duplicated. *Operations Research Letters*, *48*(4), 446-451. ("Bid request duplication is a challenging issue for display advertising buyers. […] We can picture the duplication problem this way: it is possible that several programmatic bidding agents of the same buyer receive a bid request without knowing whether the other agents have received a (duplicated) request and whether they are going to bid. Indeed, the time scale involved to answer the request is so short (less than 100 ms) that it can be technically impossible (or at least very challenging) for the buyer to synchronize the servers' behaviors.").

[367] Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -119 ("Turn off Bernanke on mediating domains on Adx").

[368] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 18 ("Disabling Bernanke on multi-calling publishers protected Google Ads and its advertisers by reducing the variance of Google Ads bids (making it less likely for Google Ads to bid more on behalf of its advertisers than would have been needed to win a potential ad opportunity in the absence of multi-calling).").

[369] *See* Presentation, "Mediation and non-second price auction: Detection and treatment" (Jun. 30, 2017), GOOG-DOJ-09875989, at -6009 ("Treatment 1: turn off Bernanke").

[370] Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -119 ("Stop buying on these mediating domains through AWBid").

for Google to detect. If Google Ads continued to bid through AwBid on impressions sold by these multi-calling publishers, it ran the risk of being called additional times through separate exchanges for the same impression. In this case, buying inventory only through AdX (with the additional protections provided by the first two treatments) protected advertisers using AwBid from overpaying for inventory as a result of duplicated calls on other exchanges.

190. These three treatments by Project Bell were designed to protect Google Ads advertisers in the face of publisher multi-calls, avoiding reductions in their online display advertising surplus and avoiding any costs that advertisers would otherwise face to adapt their bids to multi-calling on their own.[371] Google communicated these changes to publishers and encouraged them to end their usage of multi-calls.[372]

191. In a pre-launch experiment, Project Bell was estimated to increase Google Ads advertisers' conversions per dollar spent by [redacted]%.[373] This estimate likely underestimated the benefits of the program to advertisers because the experimental window was unlikely to be long enough to capture the effect of publishers reducing their use of multi-calling. A later analysis found that Google's efforts to reduce multi-calling successfully reduced the

[371] *See* Email from [redacted] to [redacted], "[Launch 188932] More aggressive treatment for mediation publishers" (Apr. 21, 2017), GOOG-DOJ-15208416, at -416 ("These launches are to protect[] advertisers from price inflation tactics by publishers and exchanges. 1) Background: Some publishers send duplicate calls to adwords in order to obtain revenue as much as possible. This causes price inflation [for] advertisers.").

[372] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 21 ("To encourage them to reduce usage of multi-calls, Google communicated with multi-calling publishers that Google Ads would be making some changes to how it submitted bids in response to multi-calling.").

[373] "Display Ads Quality Launches 2016 - Revenue OKR History" (Apr. 19, 2017), GOOG-AT-MDL-003465605, at tab "2016 Q4 Web & mApps," at cell N46.

prevalence of multi-calling on AdX by up to [redacted]%.[374] The overall effect on publishers (and Google Ads spending) was expected to be approximately neutral because spending would be redistributed from multi-calling publishers to those who did not multi-call.[375]

192. The multi-calling problem was not specific to Google. Major demand sources including The Trade Desk and MediaMath refused to work with sell-side platforms that made "duplicated requests" for the same impression.[376] [redacted]"[377] Google Ads used a milder corrective measure to protect its advertisers and other publishers without completely eliminating valuable transaction opportunities.

D. Responding to Plaintiffs' and Their Experts' Allegations about Bell

193. Plaintiffs claim that, "[i]f a publisher does not give preferential access to AdX, then Bell would drop their auctions from second- to third-price auctions."[378] This is factually incorrect in three ways.

[374] *See* "2018 Spring Perf - Tobias Maurer" (Mar. 26, 2018), GOOG-DOJ-05276941, at -944 (suggesting that Bell's impact on publisher behavior was substantial: "Launched Double call protections - together with sell-side commercialization reduced mediating pubs by 60%.").

[375] Comms Doc, "GDN Buying Change for Multiple Calls" (Jun. 7, 2017), GOOG-DOJ-10924698, at -698 ("Overall, the effect for Google is close to $0 as spend is redistributed from publishers doing multiple calls to those who don't."); Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -122 ("Note from these numbers that these domains are net consumers of pool. So, turning Bernanke off here will cause this extra pool to be automatically reinvested on non-mediating domains.").

[376] Seb Joseph, "DSPs are cracking down on bid duplication," Digiday (May 12, 2020), https://digiday.com/media/dsps-are-cracking-down-on-bid-duplication/ ("MediaMath is building a new supply chain that doesn't include SSPs that sell duplicated impressions, for example. Last month, The Trade Desk gave all the SSPs it buys impressions from two weeks to stop sending it duplicated requests to take part in the same auction.").

[377] [redacted]

[378] Fourth Amended Complaint ¶ 311.

a. *First*, Bell was a Google Ads program determining how Google Ads bid into AdX, on which Google Ads was a *bidder*: it could not, as a result, change the auction format *for AdX*.

b. *Second*, Google Ads applied Bell's treatments to publishers only when they called AdX multiple times for the same impression, not when they contracted with a non-Google exchange. If a publisher called another exchange before AdX, but did not call AdX more than once for each impression, Project Bell would not affect that publisher.[379, 380] [redacted], Senior Director of Engineering at Google, has declared that "Bell v.2 did not affect Google Ads bids on the inventory of publishers that called other exchanges before AdX, but did not call AdX multiple times, even if those other exchanges in turn called Google Ads."[381]

c. *Third*, contrary to the Plaintiffs' claim that "Bell would drop their auctions from second- to third-price auctions," when Project Bell detected a multi-calling publisher, it *disabled* Bernanke and reverted to submitting its two highest bids with a standard 14% revenue share, which *increased* Google Ads' low bid. Neither Google Ads nor AdX ever used third-price auctions.

---

[379] *See* Comms Doc, "GDN Buying Change for Multiple Calls" (Jun. 7, 2017), GOOG-DOJ-10924698, at -700 ("Q: Are publishers who work with multiple exchanges impacted by this change? A: If GDN is called multiple times in series for a single query, as the result of any of the implementations mentioned above then the [publisher's] revenue may be negatively impacted."); Design Doc, "Multiple calls and non-second price auctions: detection and management" (Sep. 23, 2016), GOOG-DOJ-AT-02471119, at -119 ("we tackle multiple calls to Adwords on just Adx").

[380] I have reviewed documents (*e.g.*, Presentation, "Mediation detection and GDN bidding" (Mar. 2, 2015), GOOG-DOJ-06563186, at -194 to -201) in which Google considered applying these treatments to publishers who used "mediated single-calls," in which publishers called another exchange before AdX. I understand these plans were never implemented. *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 16.

[381] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 16.

194. None of the Plaintiffs' experts made substantial allegations about Project Bell.[382]

[382] Professor Hochstetler provides a brief description of the functioning of Project Bell (*see* Expert Report of J. Hochstetler (Jun. 7, 2024), at Section X.E). Professor Chandler mentions "Bell" in his Summary of Opinions, but his report contains no discussion of multi-calling or the details of a program resembling Project Bell v2, and so I have interpreted his allegations to apply to Global Bernanke, also known as Project Bell v1 internally at Google (*see* Expert Report of J. Chandler (Jun. 7, 2024), at ¶¶ 23, 341).

VI. PROJECT ELMO: MANAGING ADVERTISER BUDGETS WHILE DISINCENTIVIZING BID DUPLICATION AND MULTI-CALLING

A. Overview

195. Project Elmo is a budget management feature introduced by DV360 in November 2017[383] and Google Ads in November 2018.[384] Elmo ensures that Google's buy-side tools make consistent bids on behalf of an advertiser across all bid requests received for a given end user (as identified by a cookie) within each minute.[385] By bidding consistently on behalf of an advertiser for the same end user within each minute, Elmo helps advertisers control their bidding strategy and their rates of spending—for example, to avoid rapid depletion of their advertising budgets—and disincentives the harmful practices of multi-calling and bid duplication by publishers and exchanges.

196. Plaintiffs allege that Elmo was a strategy "to reduce spend on rival exchanges[,] represent[ing] a campaign to undermine the success of header bidding and starve rival exchanges of their primary source of demand."[386] In fact, Elmo was designed to ensure

---

[383] Launch Details Spreadsheet, Launch 209956 (Aug. 29, 2023), GOOG-AT-MDL-009644201, at cells C1, C4 ("Launch Date [...] 2017-11-29").

[384] Email from [redacted] to [redacted], "[Launch 205617] Cookie-based budget throttling for GDN advertisers" (Dec. 6, 2018), GOOG-AT-MDL-015521456, at -456 ("Launch Date [...] 2018-11-19"). *See also* Email from [redacted] to [redacted], "[Launch 205617] Cookie-based budget throttling for GDN advertisers" (Dec. 6, 2018), GOOG-DOJ-AT-01363996, at -996 (noting that launch had occurred).

[385] Email from [redacted] to [redacted]m, "[Launch 201914] DBM advertiser experiment for cookie-based throttling" (Aug. 29, 2017), GOOG-DOJ-13564564, at -564 ("Cookie-based throttling means keying the throttling decision on cookie instead of making an independent decision per query [...] ."); Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -889 ("The goal is to make budget throttling consistent across the multiple queries that result from mediation. [...] To solve this, we use cookie in the throttling decision. [...] Within each one-minute window, every query for the same delivery period on the same cookie gets the same throttling decision.").

[386] Fourth Amended Complaint ¶ 405 ("Taken together, Poirot, Elmo, and other strategies to reduce spend on rival exchanges represent a campaign to undermine the success of header bidding and starve rival exchanges of their primary source of demand.").

Google's budget management tools worked correctly for its advertiser customers.[387] By blocking the harmful effects of bid duplication, Elmo reduced spending on exchanges engaged in that practice and increased spending on exchanges that did not duplicate bids, regardless of whether those exchanges participated in header bidding.[388] None of the Plaintiffs' experts conduct substantive analysis of Project Elmo; nor do they provide supporting evidence for the allegations about its adverse effects.

B. Background: Budget Throttling As a Tool To Manage Advertiser Budgets

197. As discussed in Section III.D, advertisers using Google Ads and/or DV360 select campaign parameters that determine how those buy-side tools bid on their behalf for display advertising impressions.[389] One frequently used campaign parameter is a **budget**, which limits the total spending that the buy-side tool can make on behalf of the advertiser in a designated time period. In 2017, around the time of Elmo's launch on DV360, more than [redacted]% of DV360's advertiser customers were budget-constrained, meaning that the advertiser had selected a budget and that DV360 could have purchased more impressions if the advertiser had selected a higher budget.[390]

[387] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -889 ("The goal is to make budget throttling consistent across the multiple queries that result from mediation."); Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -633 ("We addressed the following problems for DBM in 2017: [...] Budget allocation where the same query is sent to the bidder multiple times (Elmo)[.]").

[388] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -650 ("We see a significant drop in exchanges that exploited this mechanism and gains for the cleaner exchanges[.]").

[389] *See, e.g.*, Google, "Create a Display campaign," Google Ads Help (accessed Mar. 12, 2024), https://support.google.com/google-ads/answer/10759203; Google, "Create a campaign," Display & Video 360 Help (accessed Mar. 12, 2024), https://support.google.com/displayvideo/answer/7205081?hl=en.

[390] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -647 ("DBM is over [redacted]% budget constrained.").

198. Managing campaign budgets is an essential function of a buy-side tool. If a buy-side tool submits bids for *every* impression based on an advertiser's maximum willingness-to-pay *without* taking its budget into account, that advertiser would often find itself exhausting its budget quickly.[391] There are several reasons why this may be an undesirable outcome for the advertiser. *First*, many advertisers prefer to deliver ads smoothly over a campaign period, rather than very rapidly at the start of the campaign, to expand the set of users reached by the campaign.[392] *Second*, by choosing bids based on the maximum willingness-to-pay for *every* arriving impression, an advertiser might miss out on the opportunity to win later-arriving impressions at lower prices than earlier-arriving ones. *Third*, the budget management strategies of competing bidders can interact to distort pricing patterns. For example, early in the development of search advertising, Google found that cost-per-click spiked at the beginning of the day and decreased over the course of the day as bidders' daily budgets ran out, leading to an end-of-day price [redacted]% lower than at the start of the day.[393] That pricing pattern suggests a failure to manage bidder budgets well: in the absence of budget throttling, bidders would spend most of their budgets on impressions at the start of the day when competition for each impression was high, even though they could have spent the same budget to win more impressions later

[391] "Introduction to Budget Throttling" (Mar. 1, 2018), GOOG-DOJ-03792729, at -730 ("In practice there are plenty of such campaigns where the potential spend is 100x (or even more) higher than the advertiser's specified budget. In fact, lots of YouTube campaigns can exhaust their daily budget in less than [redacted].").

[392] Xu, J., Lee, K. C., Li, W., Qi, H., & Lu, Q. (2015). Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 2217-2226) ("In targeted online advertising, advertisers look for maximizing campaign performance under delivery constraint within budget schedule. Most of the advertisers typically prefer to impose the delivery constraint to spend budget smoothly over the time in order to reach a wider range of audiences and have a sustainable impact.").

[393] "Introduction to Budget Throttling" (Mar. 1, 2018), GOOG-DOJ-03792729, at -731 ("This is exactly how the system worked before probabilistic throttling was introduced in 2006, and was captured in the original probabilistic throttling design doc [...] . Basically, CPC spiked at day start, and kept decreasing, to about [redacted]% lower at day end. In fact, the problem was mitigated by budget curve (which we will discuss later in this doc) or it would have been a lot worse. Today, there is no CPC jump at day boundary").

in the day. That would lead to an inefficient outcome whenever advertisers with very high values for later-arriving impressions would fail to win them because they had already exhausted their budgets.

199. **Budget throttling** (also known as budget pacing) is a commonly-used heuristic approach by which bidders manage the rates at which their ads are shown and budgets are spent in online auctions.[394] Under budget throttling, Google would calculate bids for an advertiser for only a fraction of the impressions for which the advertiser was eligible.[395] For example, suppose that an advertiser had a budget of $100 per day to spend on a display advertising campaign, and that Google Ads—if it bid for the advertiser on *each* eligible impression without taking its budget into account—would spend $2,400 per day on eligible impressions for the advertiser. Without budget throttling (and assuming that average impression prices and numbers did not fluctuate over the course of the day), Google Ads might find that the advertiser's budget was exhausted in the first hour of the day. With budget throttling, Google Ads would instead bid on roughly one in every twenty-four impressions, leading to its budget being more smoothly spent over the course of the day.

200. Originally, Google implemented budget throttling by randomizing auction participation separately for each bid request it received.[396] In that implementation of budget throttling,

[394] Gui, G., Nair, H., & Niu, F. (2021). Auction throttling and causal inference of online advertising effects. arXiv preprint arXiv:2112.1515; Balseiro, S., Kim, A., Mahdian, M., & Mirrokni, V. (2017). Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 15-23) .

[395] "Introduction to Budget Throttling" (Mar. 1, 2018), GOOG-DOJ-03792729, at -733 ("Probabilistic throttling is the basic layer of throttling. It was first introduced in 2006. The idea is simple: it computes an impression probability (IP) for each campaign, and throttles ad candidates randomly according to IP.").

[396] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -888 ("However, budget throttling works by choosing a random number on each query.").

Google would adjust the probability of an advertiser bidding on each impression adaptively to target the advertiser's desired daily spend, increasing it if the advertiser had spent less budget than expected over the course of the day, and lowering it if the advertiser had spent more budget than expected.[397]

C. How Some Publishers and Exchanges Gamed Budget Throttling

201. By 2017, Google observed that some publishers and exchanges "exploited" Google's implementation of budget throttling to increase Google's spend on their impressions at the expense of other publishers and exchanges.[398] Because Google determined each advertiser's participation using an *independent* coin flip for each bid request, an exchange or a publisher could increase its probability of receiving a bid from a higher-value bidder on Google Ads or DV360 by sending multiple bid requests for the same impression.[399] For example, suppose DV360 had five advertisers eligible for an impression, two of whom with value $10 and three of whom with value $1 for the impression. A publisher receiving a bid of $1 from DV360 for the impression might be tempted to send one or more additional bid requests to DV360 for the same impression, hoping that DV360 would bid $10 on one of those requests. In addition to multi-calling AdX (in the way I

---

[397] *See* "Introduction to Budget Throttling" (Mar. 1, 2018), GOOG-DOJ-03792729, at -733 ("How is [Impression Probability] calculated? We use a [Proportional-Integral-Derivative] controller that computes [Impression Probability] by tracking a target spend curve.").

[398] Performance Review Document [[redacted]], "Tim perf Q3 2019 (original structure)" (Sep. 9, 2019), GOOG-DOJ-AT-02218994, at -998 ("When buying on multiple exchanges, our budget controller roughly apportioned budget based on query volume. Some exchanges exploited this by sending the same query multiple times, inflating their share of budget-constrained spend (which is most of DBM).").

[399] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -647 ("Exchanges send multiple calls for a single query to get multiple shots at budget throttling and land the largest bid [...] Money flows to publishers and exchanges that use repeated calls"); Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -888 ("However, budget throttling works by [redacted] This means that it can make different decisions for queries that actually represent the same inventory.").

discussed in Section IV), a publisher might also receive different bids from DV360 if it offered the impression via multiple exchanges. Non-Google exchanges could similarly fish for a higher bid by sending multiple bid requests for the same impression, a tactic known as **bid duplication**.[400] For example, if a non-Google exchange observed that DV360 had a high-value bidder for a certain piece of inventory around half the time it was called to bid, it might call DV360 to bid two or more times for the impression until it observed a high bid. Bid duplication has the effect of transforming the fair coin flip used to determine the advertiser represented by DV360 (and thus the bid received by the exchange) into a *series* of coin flips, repeated until the desired outcome is observed. In the words of one Google engineer, such a strategy "tricks [Google's] budget system into allocating more [...] budget to that exchange than its fair share."[401]

202. If left unaddressed, multi-calling and bid duplication can have several deleterious effects on the ad tech ecosystem.

203. *First*, these strategies undermine the budget-management techniques of buy-side tools, harming their advertiser customers. By calling a buy-side tool to bid multiple times until a desired outcome is achieved, a publisher or exchange can undo the intended effects of budget throttling, leading to bidder budgets being depleted more rapidly than bidders intend.

[400] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -647 ("Exchanges send multiple calls for a single query to get multiple shots at budget throttling and land the largest bid[.]"); Sarah Sluis, "Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication," AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication/ ("Programmatic auctions are creating so many carbon copies of themselves, it's threatening to topple the entire structure of programmatic. The bid duplication is getting so extreme, buyers are starting to take notice of this strange behavior. Instead of seeing the whole universe of bid opportunities, demand-side platforms see only a small portion of inventory copied many times, which impairs their ability to scale campaigns.").

[401] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -888.

204. *Second,* multi-calls and bid duplication raise costs for advertisers and their ad tech suppliers. The increased volume of bid requests creates engineering challenges for bidders, including attempts to "deduplicate" bid requests: a costly and imperfect process that seeks to identify when different bid requests correspond to the same impression.[402] Because the process of bid deduplication is imperfect, advertisers may miss out on valuable transaction opportunities or overpay for some impressions. Moreover, each additional call to bid creates additional time for bid requests to be processed and responded to, increasing ad latency.[403]

205. *Third*, multi-calling and bid duplication can harm publishers too. Because budgets are reallocated towards the publisher or exchange pursuing a successful multi-calling or bid duplication strategy, these strategies create a negative externality in the form of reduced payments to other publishers and exchanges.[404] This creates an incentive for the

[402] Sarah Sluis, "Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication," AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication/ ("To combat bid duplication, DSPs have turned to traffic shaping, a technique that filters excess bids using a combination of algorithms and manual selection to curate the inventory that buyers evaluate. [...] Processing billions of bid requests is expensive. In the face of skyrocketing cloud bills, SSPs and DSPs alike have implemented traffic-shaping tech as a cost-savings measure."); Sarah Sluis, "The Trade Desk Suppresses Bid Duplication Amid COVID-19 Traffic Surge," AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge ("The idea of cutting down on bid duplication predates the coronavirus pandemic. But the extra traffic has added millions in server costs, a burden shouldered largely by DSPs, as well as the agencies and marketers who pay for this cost as part of the ad tech tax.").

[403] Oded Poncz, "Traffic Duplication Might Be A Bigger Problem Than Ad Fraud," AdExchanger Opinion (Jan. 11, 2016), https://www.adexchanger.com/data-driven-thinking/traffic-duplication-might-even-be-a-bigger-problem-than-ad-fraud/ ("Another side effect of bid request duplication is that re-auctioning a bid takes time. In some cases, this could even become apparent to the end user. For example, if an end user is waiting for an interstitial ad to be shown upon opening an application, while behind the scenes there are a few auctions taking place for the same ad space, this will result in the user waiting considerably longer until they're able to access their content. This doubling of ad requests also adds to the overall latency in the ad ecosystem. If all this duplication means an ad takes five seconds to load, that makes for a poor consumer experience.").

[404] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -888 ("In particular, a single third-party exchange can send multiple queries for the same inventory, which effectively tricks our budget system into allocating more than budget to that exchange than its fair share. This trick works because the random throttling allocates budget according to the share of queries, but this allocation doesn't represent the actual

publishers and exchanges harmed by those practices to themselves adopt multi-calling and bid duplication, compounding the negative effects for advertisers and other publishers described above.[405] The responses of bidders to bid duplication, which may include reduced or less frequent bids on all impressions, can also reduce publisher revenues.[406]

D. How Elmo Protected Advertisers and Disincentivized Multi-Calling and Bid Duplication

206. DV360 launched Project Elmo in November 2017[407] and Google Ads launched it in November 2018[408] to ensure that Google's bid management algorithms worked as intended despite the onslaught of multi-calling and bid duplication tactics by exchanges and publishers.[409] Rather than randomizing participation of budget-constrained

inventory available in the case of multiple requests for the same inventory."); Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -647 ("Money flows to publishers and exchanges that use repeated calls").

[405] *See* Sarah Sluis, "Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication," AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication/ ("The issues caused by bid duplication are no secret. But, as is often the case in ad tech, being proactive is a disadvantage. Any single publisher attempting to fix the issue on their own will experience a decrease in revenue so profound that changing alone isn't an option. Meanwhile, removing waste from bid duplication could squeeze SSPs, who each get a chance to sell everything under the current setup.").

[406] Sarah Sluis, "Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication," AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication/ ("To combat bid duplication, DSPs have turned to traffic shaping, a technique that filters excess bids using a combination of algorithms and manual selection to curate the inventory that buyers evaluate. But the process doesn't actually deduplicate impressions for DSPs, and, paradoxically, the guesses made during traffic shaping can exacerbate the negative effects of duplication. [...] Traffic-shaping algorithms rely on historical data about what inventory has been bid on in the past to determine what to send buyers in the future. The danger in this approach is that buyers end up seeing an increasingly narrow view of what's out there, losing the chance to discover new inventory.").

[407] Launch Details Spreadsheet, Launch 209956 (Aug. 29, 2023), GOOG-AT-MDL-009644201, at cells C1, C4 ("Launch Date [:...] 2017-11-29").

[408] Email from [redacted] to [redacted], "[Launch 205617] Cookie-based budget throttling for GDN advertisers" (Dec. 6, 2018), GOOG-AT-MDL-015521456, at -456 ("Launch Date [...] 2018-11-19").

[409] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -888 ("Various kinds of publisher mediation (e.g. header bidding, multiple calls) can cause us to see multiple queries for the same

advertisers independently for each bid request it received, Elmo ensures that the *same* budget-constrained advertisers participate in all bid requests associated with a given cookie (see Section III.B.2) within a time period of [redacted].[410, 411] This creates consistency in the bids sent by Google Ads and DV360 across bid requests for a given cookie within the [redacted] time window. The [redacted] time window was "long enough so that multiple mediated queries are likely to fall within the same window, but short enough so that budget-constrained ad[vertisers] still have a chance to show [an ad] to any given user."[412]

207. Elmo reduces incentives for bid duplication and multi-calling designed to solicit a high bid from Google's buy-side tools. Because Elmo ensures that Google's buy-side tools choose the same sets of participants for bid requests for the same cookie that are received in close succession, an exchange contemplating a bid-duplication strategy or a publisher contemplating a multi-calling strategy can expect to receive the same bids from Google's buy-side tools on each bid request it sends, eliminating much of the benefit of such a strategy.

---

inventory. However, budget throttling works by choosing a random number on each query. This means that it can make different decisions for queries that actually represent the same inventory. In particular, a single third-party exchange can send multiple queries for the same inventory, which effectively tricks our budget system into allocating more than budget to that exchange than its fair share. This trick works because the random throttling allocates budget according to the share of queries, but this allocation doesn't represent the actual inventory available in the case of multiple requests for the same inventory. We want to change the budget throttling algorithm so that it makes a consistent decision across those queries, while still achieving the goal of hitting the impression probability overall.").

[410] Similarly, the same set of advertisers are throttled (or *not* participating) in each bid request associated with the cookie in that [redacted].

[411] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -889 ("Within each [redacted] window, every query for the same delivery period on the same cookie gets the same throttling decision."); Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -648 ("We fix the advertisers that can purchase a query from a given cookie during any specific time_bucket (budget throttling based on cookie x time bucket)").

[412] Design Document, "Cookie Budget Throttling" (Apr. 19, 2017), GOOG-DOJ-AT-02472888, at -889.

208. Meanwhile, Elmo does *not* reduce the average revenues of publishers using header bidding, Open Bidding or alternative mediation techniques to offer a single impression to multiple exchanges. Because Elmo ensures that the same advertisers using Google's buy-side tools are participating in each auction for which they are called to bid, it levels the playing field between exchanges: the winner is not merely the exchange chosen randomly to receive a high bid from Google's buy-side tools, as could happen under the previous budget throttling design. To the extent that Elmo successfully disincentivizes multi-calling and bid duplication, publishers also benefit from the elimination of the harmful externalities that they suffer when other publishers and exchanges adopt those practices.

209. A Google analysis found that Elmo decreased spending on exchanges employing bid duplication and increased it for "cleaner" exchanges not using those tactics (including AdX, United, and Improve Digital).[413] Non-Google display advertising tools (including The Trade Desk and Magnite) have also sought to combat bid duplication,[414] [redacted]

[413] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -650 ("We see a significant drop in exchanges that exploited this mechanism and gains for the cleaner exchanges [redacted].

[414] Sarah Sluis, "Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication," AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication/ ("Processing billions of bid requests is expensive. In the face of skyrocketing cloud bills, SSPs and DSPs alike have implemented traffic-shaping tech as a cost-savings measure. Magnite, for instance, bought nToggle for traffic shaping back in 2017."); Sarah Sluis, "The Trade Desk Suppresses Bid Duplication Amid COVID-19 Traffic Surge," AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge ("So two weeks ago, The Trade Desk asked exchanges to stop sending duplicate bid requests for the same ad impression.").

[redacted][415]

E. Responding to Plaintiffs' Allegations

210. Plaintiffs allege that "Google devised project Elmo to help DV360 identify when it saw the same bid request across multiple exchanges, and it decreased overall ad spend on any exchange that it suspected to meaningfully engage in header bidding,"[416] and that it "represent[ed] a campaign to undermine the success of header bidding and starve rival exchanges of their primary source of demand."[417] None of these claims are correct.

211. *First*, Elmo was designed to ensure consistent bids both when *multiple* exchanges issue bid requests for the same impression and when *a single* publisher or exchange calls Google's buy-side tools multiple times for the same impression. In this way, Elmo disincentivizes the harmful tactics of multi-calling by publishers and bid duplication by exchanges. Because of Elmo, for each impression, all exchanges are treated equally: winning the impression does not depend on being the "lucky" exchange chosen randomly by the budget-throttling algorithm to receive the high bid from Google's buy-side tools.

---

415 [redacted]

416 Fourth Amended Complaint ¶ 403.

417 Fourth Amended Complaint ¶ 405.

212. *Second*, Elmo did not penalize header bidding exchanges: it removed the advantages of bid duplication, regardless of whether the exchange participated in header bidding. Elmo did not "starve" non-Google exchanges of demand from Google's buy-side tools. Several non-Google exchanges—particularly those Google identified as not engaged in bid duplication—saw *increased* spending from Google as a consequence of Elmo.[418]

213. *Finally,* Elmo represented Google's buy-side tools *adapting* to header bidding, rather than undermining it. By protecting Google's advertiser customers from harmful tactics like bid duplication and multi-calling, Elmo ensured that its advertiser customers could safely participate in auctions for impressions on multiple exchanges, encouraging more participation. Advertisers no longer needed to fear that unscrupulous publishers and exchanges would undermine their budget-management strategies.

214. None of the Plaintiffs' experts conduct substantive analysis of Project Elmo. Professor Gans claims that "Google increased barriers to entry through Projects Poirot and Elmo," but he provides no support for that conclusion as it pertains to Project Elmo.[419] In fact, as I discussed above, Project Elmo did not increase barriers to entry; it instead treated all exchanges equally. Dr. Chandler includes Elmo in a list of programs that he claims "jeopardized, and detrimentally affected, transparency and fairness of the auctions in which they were employed."[420] He provides no analysis to support that conclusion as it pertains to Project Elmo, and his conclusion is incorrect: Project Elmo made auctions

[418] Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -650 ("We see a significant drop in exchanges that exploited this mechanism and gains for the cleaner exchanges. [redacted]).

[419] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 864.

[420] Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 23.

fairer by treating all exchanges equally, and it made participation in auctions safer for advertisers using Google's buy-side tools.

VII. POIROT AND MARPLE: IMPROVING RETURNS FOR ADVERTISERS BY OPTIMIZING BIDS IN NON-SECOND PRICE AUCTIONS

A. Overview

215. Poirot and Marple were bid optimization programs on DV360 and Google Ads, respectively, that were designed to (i) identify exchanges using non-second price auctions and (ii) bid to maximize expected advertiser surplus on those exchanges.[421] These programs benefited advertisers by increasing their returns from online display advertising, regardless of advertisers' campaign objectives. The programs made bidding easier for advertisers by automating the experiments and computations that advertisers would otherwise seek to make themselves. In doing so, they improved bidding accuracy, improving matching and reducing costs for advertisers.

216. Adopting the Plaintiffs' emphasis, I focus most of my discussion in this chapter on Poirot. Each allegation made by Plaintiffs about Poirot is either incorrect or founded on faulty assumptions:

a. Plaintiffs and their experts argue that Poirot conferred "no actual benefit to advertisers,"[422] but this is incorrect. The program both improved bidding and

[421] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -809 ("The goal of Poirot is to discover the exchanges that deviate from second pricing and bid appropriately on these to improve advertiser performance on these exchanges."), -811 ("Our optimization problem can be stated as follows: For each advertiser, find bidding policy *f(query features, ad features)* that maximizes $\mathbf{\Sigma}_i v\text{-}c_i$."); Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 26 ("In July 2017, Google launched Project Poirot, an algorithm designed to protect DV360 advertisers from overbidding on exchanges that deviated from second-price auctions"). *See also* Presentation, "Poirot Launch Metrics" (Oct. 5, 2021), GOOG-DOJ-AT-02480338, at -341; Email from [redacted] to [redacted] "[Launch 187971] Project Poirot (**DRAFT**)" (Apr. 11, 2017), GOOG-DOJ-14398809, at -810 ("This launch changes bids for fixed CPM DBM advertisers to maximize advertiser surplus."); Presentation, "DV360 optimizations ENG deep dive" (Jan. 24, 2020), GOOG-DOJ-11733552, at -577.

[422] Fourth Amended Complaint ¶ 400 ("[Poirot was] a direct reallocation of advertising dollars to Google's own ad exchange with no actual benefit to advertisers."). *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865 ("Poirot

increased advertiser surplus, as verified by Google's internal experiments.[423] Poirot also protected DV360 advertisers from misleading tactics from exchanges that ran so-called "dirty auctions," in which the exchange claimed to use one auction format to sell an impression, but actually used another.[424]

b. Plaintiffs and their experts allege that Poirot was designed to shift advertiser spending from header bidding to AdX, but this, too, is incorrect.[425] Poirot applied to non-second-price auctions, regardless of whether the exchange participated in header bidding. It did not shade bids into exchanges using second-price auction rules,[426] nor did it shade bids below the value that maximized advertiser surplus.[427]

resulted in reallocating revenue from rival exchanges to Google's own exchange, which had no benefit to advertisers or publishers.").

[423] *See* Email from [redacted] to [redacted] et al., "Re: Poirot to launch 6/19" (Aug. 20, 2017), GOOG-DOJ-07825115, at -115 ("Through experiments, we measured that […] the surplus increases by 12% in the affected exchanges as a result of this launch."); Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -815 ("The following table shows the impact limited to the non-second price auction exchanges."; "surplus [...] change[:] 15.57%"). Google experiments found that Poirot increased advertiser surplus by 6% on all exchanges, including second-price exchanges. *See* Presentation, "Bidding in Adversarial Auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Advertiser impact […] 6% surplus increase").

[424] Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -635 to -636 ("There are roughly three types of auctions[:] Second price (buyer bids truthfully)[,] First price (buyer has to shave bids)[,] Dirty (called second price, but really more like first price) […] Project Poirot ensures advertiser bids are protected").

[425] Fourth Amended Complaint, Section VII.D.3.v ("Google diverts ad spend away from rival exchanges that engage in header bidding."); Expert Report of J. Gans (Jun. 7, 2024), at ¶¶ 863-64 ("In addition to the conduct that I have found to be anticompetitive in itself, there were additional actions aimed at limiting the impact of Header Bidding [...] As one example, Google increased barriers to entry through Projects Poirot and Elmo.").

[426] Such exchanges include United and Improve Digital. *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 33 ("Poirot also found that reducing bids into some other ad exchanges (such as Improve Digital and United) did not increase advertiser surplus by more than the [redacted] percent threshold, so Poirot did not lower DV360 bids into those exchanges either.").

[427] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 31 ("Poirot's process for calculating multipliers that maximized expected advertiser surplus did not take into account whether an ad exchange participated in header bidding."); Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -811 ("Our optimization problem can be stated as follows: [redacted]

c. Plaintiffs also allege that Poirot redirected ad spend to AdX when it "engaged in exactly the same behavior" that DV360 criticized, but from September 2017 (just two months after its launch in July 2017) up until the transition to the UFPA in September 2019, Poirot applied in just the same ways to AdX as to other exchanges.[428, 429] After the transition to the UFPA, DV360 optimized bids to the new AdX format with an updated Poirot algorithm that used the minimum-bid-to-win information provided by AdX to all bidders.[430]

d. Plaintiffs characterize Poirot as a "strateg[y] to reduce advertiser spend on rival exchanges,"[431] but this analysis relies on two false assumptions: that advertisers would not otherwise shade bids on their own and that advertiser savings from Poirot would go unspent.

[428] Fourth Amended Complaint ¶ 400 ("[With Poirot,] DV360 was actually redirecting that ad spend to a marketplace that engaged in exactly the same behavior.").

[429] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶¶ 33, 35-36 ("Since September 2017, Poirot has applied on AdX and on all third-party ad exchanges on which DV360 bids."); Deposition of [redacted] at 274:9-13 (Sep. 17, 2021), GOOG-AT-MDL-007173084, at -358 ("A. Google's buy side algorithms run across all the traffic. So that—whether it's AdX or third-party exchanges, the algorithm that we are talking about [Poirot] applies on all the traffic for DV360.").

[430] The minimum-bid-to-win information is provided to bidders at the conclusion of an auction for an impression on AdX and tells the bidder the minimum bid it would have needed to make to win an auction. *See* Google, "Bid data sharing," Authorized Buyers Help (accessed Jul.Dec. 24, 2024), https://support.google.com/authorizedbuyers/answer/2696468?hl=en ("[W]hen a bidder submits a valid bid into the auction, they receive back the minimum value they would have had to bid to win that auction, whether they lost or won."). *See also* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 35 ("Poirot was also updated following Google's shift to a Unified First Price Auction. With the transition to a Unified First Price Auction, Google began providing minimum-bid-to-win data to buyers, and DV360 began to use that minimum-bid-to-win data to inform how Poirot would lower bids into AdX in order to optimize for expected advertiser surplus.").

[431] Fourth Amended Complaint ¶ 405 ("Taken together, Poirot, Elmo, and other strategies to reduce spend on rival exchanges represent a campaign to undermine the success of header bidding and starve rival exchanges of their primary source of demand."). *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865 ("Poirot resulted in reallocating revenue from rival exchanges to Google's own exchange, which had no benefit to advertisers or publishers. They are examples of the actions that Google could take to alter the flow of information into markets in ways that were not motivated by the needs of Google consumers (publishers/advertisers) but could disrupt the efficient operation of markets in ways that potentially reduced match quality and, potentially, and made it more difficult for rival exchanges and entrants to compete.").

B. Background: Non-Second-Price Auctions

217. As I discussed in Section III.C.4, many ad exchanges switched from using second-price auctions to non-second-price auctions for online display advertising impressions between 2017 and 2019.[432] During this transitional period, different exchanges used different auction formats, and some exchanges even used different auction formats for different impressions.[433] Some exchanges ran what Google engineers called **dirty auctions**, meaning that they claimed to run a second-price auction but actually tried to extract additional payments from bidders by charging the winner a price between its own bid and the highest losing bid.[434] In this section, I refer to both first-price auctions and dirty auctions as **non-second-price auctions**.

218. To illustrate a dirty auction, suppose an exchange sets a floor price of $4.00, and then observes that the only bid above the floor price is $5.00. In a second-price auction, the bidder would pay $4.00. In a dirty auction, the exchange might charge a higher price, such as $4.60, claiming that was the second-highest bid. From the price and its bid alone, a bidder would be unable to detect that the exchange had charged a price different from the second-highest bid.

[432] *See* Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -267 (Figure: "Impression Share (US Ad Inventory)," "First-price auction": [redacted] in December 2017, [redacted]% in March 2018, "Second-price auction with anomalies": [redacted]% in December 2017, [redacted]% in March 2018).

[433] *See* Presentation, "Poirot with auction type signal" (Nov. 5, 2018), GOOG-DOJ-05283173, at -185 (Figure: "Impression volume per auction type").

[434] Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -635 ("There are roughly three types of auctions[:] Second price (buyer bids truthfully)[,] First price (buyer has to shave bids)[,] Dirty (called second price, but really more like first price)").

219. As I have emphasized (see Section III.C above), a surplus-maximizing bidder must bid differently depending on the auction's pricing rule.[435] This created new challenges for bidders in the transitional period in which exchanges differed in their pricing rules: how to predict the auction format used for a given impression, and how to optimize bids to that format. As I have discussed above, second-price auctions are bidder-truthful, meaning that bidders maximize their surplus by bidding their values, whereas first-price auctions are not, so that bidders in these auctions maximize surplus by shading their bids. Dirty auctions also fail to be bidder-truthful and may be even more challenging for bidders, since a bidder may need to monitor its auction performance over time to detect changes in the auction format and adapt bids accordingly. The benefits of bid-shading were widely understood in the online display advertising industry.[436]

C. All Advertisers Should Maximize Expected Surplus

220. Online display advertising campaigns can have many different objectives, which may involve acquiring a fixed number of impressions in some time period or maximizing clicks subject to a fixed budget or allocating a budget between advertising on the web or other media, and others. Regardless of the advertiser's campaign objective, it can always

[435] *See, e.g.*, Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 110-16.

[436] S*ee, e.g.*, Jessica Davies, "What to know about Google's implementation of first-price ad auctions," Digiday (Sep. 6, 2019), https://digiday.com/media/buyers-welcome-auction-standardization-as-google-finally-goes-all-in-on-first-price/ ("[D]emand-side platforms came up with bid shading as a way to help buyers transition to first-price auctions where they have to be willing to pay what they bid"); Digiday, "In programmatic, buyers sometimes don't know what type of auction they're bidding in," Digiday (Jun. 30, 2017), https://digiday.com/marketing/ad-buyers-programmatic-auction/ ("A buyer might think they're buying based on second price but really be in a first-price auction. That can get expensive, since the bid strategies are far different"); Sarah Sluis, "Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price," AdExchanger (Sep. 5, 2017), https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/ ("[I]f programmatic traders think they are still playing according to second-price auction rules, they will overpay for inventory. To combat price increases, some buyers have already started shading, or reducing bid prices.").

acquire the same number of impressions at *minimum* cost by adopting a bidding strategy that *maximizes* the expected surplus from each auction opportunity, using a properly determined value for each individual impression.[437] A surplus maximization strategy (which is the type of strategy implemented by Poirot) can thus serve the interests of every advertiser, regardless of its campaign objectives, by achieving those objectives at minimum cost.

D. Poirot Benefited Advertisers By Optimizing Bids to Auction Formats

1. Overview of Project Poirot

221. Poirot was a DV360 program to make it easier for advertisers on DV360 who set up their campaigns using **fixed CPM** bidding (also known as **manual bidding**) to implement optimal bidding strategies.[438] Fixed CPM bidding is an option in DV360 that allows advertisers to specify an amount that DV360 "can spend to win any individual impression

[437] This claim relies on the two assumptions that (1) the bidder can win any fraction of similar auctions for an impression by varying its bid from low to high and (2) the bidder (or its DSP) accurately estimates the distributions of bids in the auctions. Let us begin with some arbitrary bidding strategy *B* for the advertiser. Assumption (1) implies that there is some value that it can use with expected-value-maximization bidding that wins the same expected number of impressions as bidding strategy *B*. Assumption (2) implies that the *expected* number of impressions and *expected* surplus are equal to the *realized* numbers. Since both strategy *B* and the new strategy win impressions of the same total value and the expected-surplus-maximizing strategy has a higher surplus, it must have a lower cost.

[438] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -811 ("Note that this project is only applied to Fixed CPM bidders."). *See also* Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -285 ("To date, fixed bidding in DBM has always bid the the [*sic*] exact inputted CPM/1000 for every impression. Now, in order to ensure advertisers are getting the best possible price for each impression, we are preparing to launch an optimization with the goal of winning the same impression for a lower price. For Optimized Fixed CPM Bidding [Poirot], the inputted CPM value will serve as a maximum CPM bid, as opposed to a fixed CPM bid - we will never bid more than the inputted value.").

for a line item."[439] In 2017, most DV360 advertisers used manual bidding, although automated bidding strategies have since become more popular.[440]

222. Before Poirot, fixed CPM bidding used the same bid for each impression, regardless of the auction format being used for a given impression.[441, 442] A built-in functionality to adapt bids to the auction format was not demanded when the vast majority of exchanges used second-price auctions, but its absence became a problem when exchanges started to use different pricing rules because advertisers needed to adapt their bids to the auction format in order to maximize their profits. Project Poirot launched fully in July 2017[443] and replaced fixed CPM bidding with **optimized fixed CPM bidding**, in which DV360 treated the CPM reported by a DV360 advertiser as its value (its optimal bid into a clean second-price auction) and optimized the advertiser's bid to the auction format of the exchange offering the impression.[444]

---

[439] Google, "Set a fixed CPM bid for a line item," Display & Video 360 Help (accessed Jul.. 24, 2024), https://support.google.com/displayvideo/answer/2696858?hl=en.

[440] *See* Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -265 ("Auto-bidding Adoption[:] Jan 2017: 17% [,] Sep 2018: 40%").

[441] *See* Presentation, "DV360 optimizations ENG deep dive" (Jan. 24, 2020), GOOG-DOJ-11733552, at -553 ("DV360 three years ago: Mostly fixed CPM manual bidding").

[442] Advertisers could set up separate line items for different exchanges (assuming they knew the auction formats), but at the time of Poirot's launch, "barely any" advertisers set exchange-specific fixed CPMs. *See* Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -286. Even so, exchange-specific line items would be insufficient to allow advertisers to optimize bids to the "auction type" signals as in Poirot v2 (see Paragraph 235 below).

[443] *See* Email from [redacted] to [redacted] and [redacted], "Metrics post Poirot launch" (Jul. 24, 2017), GOOG-DOJ-05270417, at -417 ("Poirot was launched fully on July 19, 2017.").

[444] *See* Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -285 for an explanation of fixed CPM bidding and optimized fixed CPM bidding (Poirot) ("To date, fixed bidding in DBM has always bid the the [*sic*] exact inputted CPM/1000 for every impression. Now, in order to ensure advertisers are getting the best possible price for each impression, we are preparing to launch an optimization with the goal of winning the same impression for a lower price. For Optimized Fixed CPM Bidding [Poirot], the inputted CPM value will serve as a maximum CPM bid, as opposed to a fixed CPM bid - we will never bid more than the inputted value.").

223. Aptly named after a famous fictional sleuth, Poirot was designed to "discover the exchanges that deviate from second pricing and bid appropriately on these to improve advertiser performance on these exchanges."[445] To discover the exchanges deviating from second-price rules, Poirot conducted experiments to determine whether bidding an advertiser's value on an exchange was an optimal or near-optimal strategy. I describe these experiments in detail below. Poirot then adjusted advertisers' bids into the exchanges that it found were deviating significantly from second-pricing in order to maximize the expected advertiser surplus.[446] Poirot applied the same experiment-and-optimize algorithm to all non-Google exchanges from its launch in July 2017, and it applied the same algorithm to AdX from September 2017 onwards.[447]

224. Although DV360 expected that most fixed CPM advertisers would benefit from Poirot, it notified advertisers that they could "opt out" of optimized fixed CPM bidding by unchecking a box in the DV360 user interface.[448] Fewer than [redacted]% of affected advertisers chose to opt out.[449]

[445] Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -809.

[446] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -636 to -637; Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -811 ("Our optimization problem can be stated as follows: For each advertiser, find bidding policy [redacted]

[447] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶¶ 33, 36 ("Before Google transitioned to a Unified First Price Auction, Poirot determined that reducing bids into AdX did not increase expected advertiser surplus by more than the 10-percent threshold, so Poirot did not lower DV360 bids into AdX. […] Since September 2017, Poirot has applied on AdX and on all third-party ad exchanges on which DV360 bids."); Deposition of N. Jayaram at 274:9-13 (Sep. 17, 2021), GOOG-AT-MDL-007173084, at -358 ("A. Google's buy side algorithms run across all the traffic. So that—whether it's AdX or third-party exchanges, the algorithm that we are talking about [Poirot] applies on all the traffic for DV360.").

[448] *See* Design Doc, "Project Poirot" (Mar. 31, 2017), GOOG-DOJ-11247631, at -631 ("We're proposing to add a checkbox at partner level with per-advertiser option to opt-out"); Email from [redacted] to [redacted], "Re: [Update] Project Poirot stats" (May 3, 2017), GOOG-DOJ-12025827, at -827 ("As part of step 2, we'll communicate to advertisers and they'll have the option to opt out. This will happen before full launch, which is step 3.").

[449] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Very few customers [redacted]%) opted out.").

2. How Poirot Experimented to Identify Non-Second-Price Auctions and Optimize Bids

225. Poirot used experiments to determine non-second-price auctions and optimize bids. The first version of Poirot performed this experiment-and-optimize process separately for each advertiser and each exchange.[450] Every day, for a subset of the advertiser's traffic to each exchange, Poirot tested bids between [redacted]% and [redacted]% of the advertiser's reported CPM and estimated the advertiser surplus for each multiplier that it tested in that range.[451] It then fit a curve to the results of these experiments, such as the one shown in Figure 6, and identified the bid multiplier that maximized the advertiser's expected surplus.[452] If the exchange was using a second-price auction, economic theory predicts that the multiplier maximizing the advertiser's expected surplus should be 1 (resulting in bids equal to advertiser values), reflecting the fact that the second-price auction incentivizes a bidder to bid its value for an impression. If Poirot found that a different multiplier significantly increased the advertiser's expected surplus, that would suggest that the exchange was likely to be using a non-second-price auction.

[450] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -812 ("Per-exchange is the minimum requirement […] Customer id is [also] used"); Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -279 ("Surplus Change vs. Bid Multiplier[:] Measured for each advertiser x exchange, from daily background experiment"). This process was performed unless an advertiser had insufficient traffic to a given exchange to conduct these experiments, as discussed in Paragraph 226.

[451] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -813 ("We run a nightly data pipeline that gathers the necessary metrics (original bid, cost, impressions) for the relevant slices and use the previous [redacted] days to optimize."); Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -279 ("Surplus Change vs. Bid Multiplier[:] Measured for each advertiser x exchange, from daily background experiment"); Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -682 ("Current Poirot Version […] Methodology: [redacted]

[452] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -813 [redacted]; Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -279 ("Surplus Change vs. Bid Multiplier[:] [redacted]

Figure 6: Poirot multipliers and advertiser surplus.[453]

226. Because this process involved daily experiments for each advertiser on each of several exchanges, the number of experiments created the likelihood of **false positives** in which random statistical noise might lead the test to wrongly conclude that a second-price exchange was using a non-second price format. Poirot included two protections to reduce the likelihood of this error. *First,* if an advertiser had insufficient traffic to a given exchange to conduct the Poirot experiments with sufficient statistical confidence, DV360 used a default multiplier for that exchange, calculated using an experiment conducted on DV360's overall traffic to the exchange.[454] As I discuss in Paragraph 235 below, a later update to Poirot removed advertiser-specific optimizations altogether, opting instead to use data collected at the exchange level. *Second*, Poirot adjusted bids only if the

[453] Presentation, "Poirot with auction type signal" (Nov. 5, 2018), GOOG-DOJ-05283173, at -175.

[454] *See* Presentation, "Poirot Review" (Jun. 10, 2019), GOOG-DOJ-32261273, at -289 to -290 ("Exchange priors (used if insufficient data for given advertiser)").

experiment found that the optimal multiplier would increase advertiser surplus by at least [redacted]%. Google set the [redacted]% threshold to "avoid changing bids on second-price auctions due to noise in the data" because engineers "want[ed] to be confident [that] an exchange is unclean before lowering bids on it."[455] Based on this threshold, the Poirot algorithm typically chose to submit bids unadjusted into exchanges it predicted to be running second-price auctions, including AdX (then using a second-price auction format) and the Improve Digital and United ad exchanges (and potentially other exchanges, as well).[456]

227. For concreteness, suppose that DV360 experiments on a small subset of an advertiser's bids into an exchange, adjusting the advertiser's bids using bid multipliers between 0.6 and 1, and obtains the data on average advertiser surplus displayed in Table 2 below.

Table 2: Example Data on Average Advertiser Surplus

Bid Multiplier	0.6	0.7	0.8	0.9	1
Average Advertiser Surplus per thousand impressions	$0.30	$0.35	$0.35	$0.30	$0.20
Surplus Relative to Bid Multiplier 1	1.5	1.75	1.75	1.5	1

228. In this data, bid multipliers less than 1, corresponding to bids less than the advertisers' values, improve advertiser surplus. This suggests that this exchange was unlikely to be

[455] Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -813 ("If the maximum surplus change is less than [redacted]%, we just select [redacted]. This is to avoid changing bids on second-price auctions due to noise in the data; we want to be confident than an exchange is unclean before lowering bids on it."). *See also* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 30 ("If a multiplier increased the expected advertiser surplus by less than 10 percent, then Poirot would not adjust the advertiser's bids into that ad exchange. The purpose of this [redacted]-percent threshold was to avoid adjusting bids on second-price auctions that might erroneously appear to be non-second-price auctions due to noise in the data.").

[456] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -641 to -642.

using a clean second-price auction format because clean second-price auctions have optimal bid multipliers of 1.

229. Under Poirot v1, DV360 would calculate the best-fitting quadratic curve for this data, as illustrated in Figure 7. It would then identify the bid multiplier that maximizes the relative surplus using that curve, which is 0.75 for this data. Poirot would then multiply the advertiser's subsequent bids into that exchange on that day by 0.75.

Figure 7: Example calculation of Poirot bid multipliers[457]

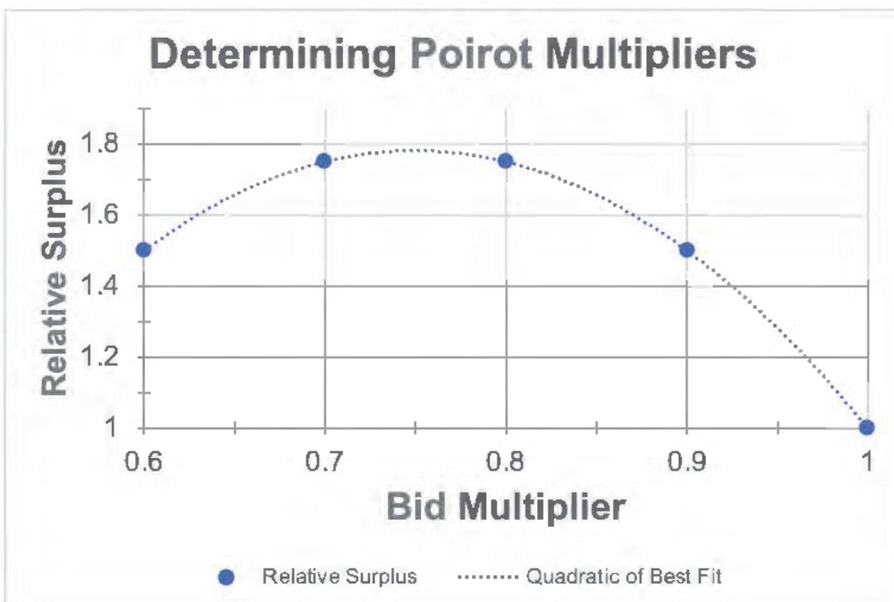

3. Poirot Increased Advertiser Surplus

230. Google studies confirmed that this first version of Poirot benefited advertisers. Two internal studies estimated that Poirot increased advertiser surplus on non-second-price

[457] This figure was created from [redacted] in my supporting materials. The figure is saved in [redacted]

exchanges by [redacted]%.[458] A similar internal study found that Poirot increased advertiser surplus by [redacted]% in total for all exchanges, including second-price exchanges.[459]

231. Poirot shaded bids only enough to maximize advertiser surplus. If DV360's goal was to reduce spending on competing exchanges, it could have shaded bids beyond the point that Poirot determined to be optimal, but it did not do that. Moreover, the Poirot procedure did not apply differently to exchanges that participated in header bidding.[460] These observations are inconsistent with the interpretation that Poirot was designed to undermine header bidding.

232. Poirot protected advertisers and increased advertiser surplus in two ways. First, its direct effect was to reduce the price of impressions purchased on non-second-price exchanges, including those "dirty" exchanges that were unclear or untruthful in describing their auction formats. Second, the budget saved on those impressions allowed budget-constrained advertisers to bid on and win *more* impressions than previously.[461] Total spending increased on AdX and some non-Google exchanges (including United and

---

[458] *See* Email from [redacted] to [redacted] et al., "Re: Poirot to launch 6/19" (Aug. 20, 2017), GOOG-DOJ-07825115, at -115 ("Through experiments, we measured that […] the surplus increases by [redacted]% in the affected exchanges as a result of this launch."); Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -815 ("The following table shows the impact limited to the non-second price auction exchanges."; "surplus […] change[:] [redacted]%").

[459] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Advertiser impacts […] [redacted]% surplus increase ($252M)").

[460] Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 31 ("Poirot's process for calculating multipliers that maximized advertiser surplus did not take into account whether an ad exchange participated in header bidding.").

[461] There may have even been another indirect effect, which is that advertisers may have increased their budgets in response to better value per dollar spent enabled by Poirot.

Improve Digital) but decreased on others.[462] Overall, Google's analysis found that Poirot did not affect advertisers' *total* spend, suggesting that budgets were redistributed.[463] The fact that only [redacted]% of advertisers opted out of Poirot also speaks to advertisers' perceptions of its benefits.[464]

4. How DV360 Updated Poirot to Further Increase Benefits to Advertisers

233. Over time, there were three major updates to the Poirot program.

234. The first, internally called **Poirot with Bid Buckets**, launched in January 2018.[465] This update grouped an advertiser's bids for different impressions into five ranges, known as "bid buckets," depending on the amount of the bid, and calculated different bid multipliers for each (using the same experimentation process described above).[466] This approach helped to implement optimal bid shading in non-second-price auctions, in which auction theory observes that the profit-maximizing bid multiplier can vary depending on the bidder's value. Google's pre-launch experiments anticipated that Poirot

---

[462] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Exchange impact[:] Overall spend neutral[;] Spend and CPM on dirty auction exchanges dropped by [redacted]%[;] Spend up by [redacted]% on second price auction exchanges").

[463] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Exchange impact[:] Overall spend neutral"); Email from [redacted] to [redacted] et al., "Re: Poirot to launch 6/19" (Aug. 20, 2017), GOOG-DOJ-07825115, at -115 ("Overall DBM spend is neutral, indicating no difficulty in redistributing the budgets. Budget reports also do not show any increase in under-delivery.").

[464] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Very few customers ([redacted]%) opted out.").

[465] Email from [redacted] to [redacted], "OVERDUE LAUNCH - Please update: [Launch 215784] Poirot: Bid bucket surplus model" (Jan. 10, 2018), GOOG-DOJ-13579782, at -782 ("Launch Date[:] 2018-01-08").

[466] *See* Design Doc, "Summary of Poirot with Bid Buckets" (Jan. 2018), GOOG-DOJ-13619370, at -370-371 ("In this launch, we are adding 1 extra feature to the Poirot model which predicts advertiser surplus as a function of bidding change. This feature is bid_bucket, and it helps capture non-second-pricing from soft floors.").

with Bid Buckets would improve advertiser surplus for fixed CPM traffic by an additional [redacted]% to [redacted]%.[467]

235. The second major update was called **Poirot v2.0**, which launched in September 2018[468] and made four changes to the Poirot algorithm:

a. [redacted] This fourth change

467 *See* Email from [redacted] to [redacted]," OVERDUE LAUNCH - Please update: [Launch 215784] Poirot: Bid bucket surplus model" (Jan. 10, 2018), GOOG-DOJ-13579782, at -783 ("From 1% experiments: […] Over fixed CPM DBM traffic on all exchanges Surplus [redacted]% […] >From 5% experiment […] Over fixed CPM DBM traffic on all exchanges Surplus [redacted]%."); Design Doc, "Summary of Poirot with Bid Buckets" (Jan. 2018), GOOG-DOJ-13619370, at -371 ("Surplus expt/cntl […] Fixed CPM DBM x external exchanges [redacted]").

468 *See* Launch Details Spreadsheet, Launch 259738 (Aug. 25, 2023), GOOG-AT-MDL-009644238, at cells C1, C2 ("Launch Date [...] 2018-9-6").

469 Some exchanges would vary their format on an impression-by-impression basis and report the format to bidders using an "auction type" signal. *See* Presentation, "Poirot with auction type signal" (Nov. 5, 2018), GOOG-DOJ-05283173, at -176 ("Since January 2018, 3PE have started providing the type of auction they are running using bid_request").

470 *See* Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -682 to -683 ("Lower the floor on bid shaving from [redacted] […] Change the model from [redacted] to improve fits and accuracy. […] Remove customer id. Past experiments have shown surplus gains when customer id is excluded from the model").

was adopted after Google engineers determined that [redacted].[471]

Pre-launch experiments of Poirot v2.0 anticipated an increase of surplus for fixed CPM advertisers on non-Google exchanges of [redacted]%.[472] By reducing the prices of impressions won on non-second-price exchanges, Poirot v2.0 allowed advertisers with limited budgets to purchase more impressions.

236. Poirot's third major update was in 2019 after AdX's transition to a unified first-price auction. This update [redacted][475]

---

[471] *See* Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -683 ("Remove customer id. Past experiments have shown surplus gains when customer id is excluded from the model").

[472] *See* Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -682 ("On 3PE (third party exchanges) we see an aggregate surplus increase of [redacted]% over all DBM traffic and [redacted]% over Fixed CPM DBM traffic.").

[473] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 35 ("Poirot was also updated following Google's shift to a Unified First Price Auction. With the transition to a Unified First Price Auction, Google began providing minimum-bid-to-win data to buyers, and DV360 [redacted].

[474] *See* Design Doc, "DBM HDMI Consolidated" (May 11, 2020), GOOG-AT-MDL-002293467, at -468 ([redacted]

[475] *See* Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 35 ("Google subsequently updated Poirot to use models built from Google's minimum-bid-to-win data to optimize for [redacted])

237. In September 2018, Google launched Marple,[476] a program on Google Ads that optimized advertiser bids in a similar manner to Poirot. Google Ads (through AWBid) could bid on behalf of advertisers for certain inventory on non-Google exchanges. Marple, initially called "Poirot for AWBid,"[477] performed a similar experiment-and-optimize procedure to increase advertiser surplus for Google Ads advertisers.[478] The initial version of Marple chose bid multipliers to maximize advertiser surplus for each bid bucket on each exchange.[479]

E. Responding to Plaintiffs' Allegations

1. Poirot Benefited Advertisers

238. Plaintiffs and their experts incorrectly assert that Poirot conferred "no actual benefit to advertisers."[480] This allegation is incorrect: Poirot benefited advertisers by protecting them from overbidding into "dirty" auctions and other non-second-price auctions. By

[476] *See* Launch Details Spreadsheet, Launch 258064 (Aug. 25, 2023), GOOG-AT-MDL-009644236, at cells C1, C3 ("Launch Date[…] 2018-9-10").

[477] Design Doc, "Poirot for AWBid Design Doc" (Sep. 10, 2018), GOOG-DOJ-AT-02512863, at -863 ("Motivated by project Poirot, project Poirot for AWBid (A.K.A Marple) aims to provide optimal bidding strategy for GDN advertisers to buy on non-second price exchanges.").

[478] *See* Design Doc, "Poirot for AWBid Design Doc" (Sep. 10, 2018), GOOG-DOJ-AT-02512863, at -864 ("As the first attempt to develop an optimal strategy to adjust the bids for AWBid, we borrow the success of Poirot project and start with the current modeling approach […] of Poirot in production."); Email from [redacted] "[Launch 258064] Project Marple" (Aug. 14, 2018), GOOG-DOJ-15264552, at -552 ("We always knew that some external exchanges deviate from second pricing. We developed an algorithmic framework to detect and quantify this deviation using AdWords data. Using this framework, we have built bid optimizations to protect AdWords advertisers against price gouging in these 'unclean' exchanges. In this launch, in response to exchanges running non-second price auctions, we lower bids for AdWords buyers bidding into these exchanges in an algorithmic fashion. The algorithm aims to find the optimal bid decrease that maximizes advertiser surplus.").

[479] *See* Design Doc, "Poirot for AWBid Design Doc" (Sep. 10, 2018), GOOG-DOJ-AT-02512863, at -864 to -865 ("We then adjust the advertiser's bid in order to maximize the [redacted]

[480] Fourth Amended Complaint ¶ 400. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865 ("Poirot resulted in reallocating revenue from rival exchanges to Google's own exchange, which had no benefit to advertisers or publishers.").

choosing bids that maximized advertisers' expected surplus on each impression, Poirot allowed advertisers to pursue *any* campaign objective at the minimum cost. Poirot improved DV360 as a product, performing a bid optimization procedure that advertisers would otherwise want to pursue themselves. Advertisers were given the option to opt out of Poirot, and fewer than [redacted]% of advertisers elected to do so.[481]

239. Plaintiffs allege that "Poirot would typically adjust DV360's bid to avoid [...] providing the rival exchange with meaningful data about DV360's willingness to pay."[482] This, however, is a side-effect of optimal bidding, which competing exchanges should anticipate. If an advertiser ever bids the same amount for an impression in a non-second-price auction as it does for an identical impression in a second-price auction, then it is necessarily a mistake: there is always a way to change its bids to buy the same number of impressions at lower cost. This means that failing to detect a non-second-price auction and to adapt bids is also a mistake.

240. I illustrate how this works for first- and second-price auctions with an example, for which the arithmetic is displayed in Table 3. Suppose some advertiser is bidding $0.50 per thousand impressions in a first-price auction and winning 1.02 million impressions, and that advertiser is bidding the same amount in a second-price auction, winning some other number of similar impressions. Suppose the bidder reduces its bids to win 20,000 fewer impressions in the first-price auction and increases its bids to win 20,000 more in the second-price auction, leaving the total number of impressions it wins unchanged. For

[481] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Very few customers ([redacted]%) opted out.").

[482] Fourth Amended Complaint ¶ 400.

simplicity, suppose it accomplishes this by reducing its bid in the first-price auction by, say, \$0.01 while increasing its bid in the second-price auction by, say, \$0.04. On the one million "unchanged" impressions that it continues to win in the first-price auction, but at a lower price, it saves \$0.01 per thousand impressions, or \$10.00 in total, as shown in Table 3. On the unchanged impressions that it continues to win in the second-price auction with its new higher bid, its prices do not rise despite its higher bid, because prices paid for any impression won in a second-price auction *do not depend on the winner's bid*. On the 20,000 "switched" impressions for which it formerly paid \$0.50 per thousand in the first-price auction, it now pays prices between \$0.50 and \$0.54, or about \$0.52 on average—\$0.02 more than it would have previously paid for those impressions. In total, the advertiser pays about \$0.40=\$0.02×20 more than before for the switched impressions, so its net savings is approximately \$9.60.

Table 3: Example Bidding in First-Price Auctions

	A	B	C=AxB
Impression Category	**Impressions** *(in thousands)*	**Savings (loss)** *(per thousand)*	**Money Saved (lost)**
Unchanged (1st price)	1,000	.01	\$10.00
Unchanged (2nd price)	Any number	0	0
Switched	20	(.02)	(\$0.40)
Total			\$9.60

241. Although the numbers in the example are specific, the conclusions are general: the advertiser saves money by reducing its bid into the non-second-price auction and increasing its bid into the second-price auction while winning the same number of impressions. The money saved on the unchanged impressions in the first-price auction is proportional to the number of those impressions, while the additional cost of the switched

impressions is proportional to the number switched. Because the number of unchanged impressions in the first-price auction is so much larger than the number of switched impressions, the advertiser is better off with the revised bids.

242. Plaintiffs also suggest that Poirot was created "ostensibly" to "avoid optimizations that were bad for advertisers."[483] Professor Gans expands this to the remarkable claim that the reallocation of spending under Project Poirot "had no benefit to advertisers."[484] Yet confirming the theory described above, Google's internal studies estimated that Poirot would increase advertiser surplus on impressions purchased through non-second-price exchanges by [redacted]%.[485] As I also described above, Poirot further aided advertisers by performing bidding optimizations that advertisers would otherwise seek to perform themselves. In suggesting that increasing advertiser surplus was merely a pretext for launching Poirot, the Plaintiffs' narratives ignore Poirot's sizable benefits for advertisers.

243. Professor Gans also suggests that Poirot "potentially reduced match quality,"[486] but I see no reason to expect that outcome in either theory or practice. By ensuring that Google Ads advertisers did not overpay for impressions (also allowing budget-constrained advertisers to spend those savings on additional impressions), Poirot would most likely have improved match quality, not reduced it. This is because bidders can use savings

[483] Fourth Amended Complaint ¶ 400.

[484] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865.

[485] *See* Email from [redacted] to [redacted] et al., "Re: Poirot to launch 6/19" (Aug. 20, 2017), GOOG-DOJ-07825115, at -115 ("Through experiments, we measured that […] the surplus increases by [redacted]% in the affected exchanges as a result of this launch."); Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -815 ("The following table shows the impact limited to the non-second price auction exchanges"; "surplus […] change[:] [redacted]%"). Google experiments found that Poirot increased advertiser surplus by [redacted]% on all exchanges, including second-price exchanges. *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -644 ("Advertiser impact […] 6% surplus increase").

[486] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865.

from Poirot to compete for more impressions, increasing thickness and likely increasing matching quality.

2. Poirot Targeted Non-Second-Price Auctions, Not Header Bidding As Alleged

244. Plaintiffs and their experts allege that Poirot was a scheme to "combat" header bidding,[487] but this inference is inconsistent with the design of Poirot in at least three ways.

245. *First*, Poirot lowered bids to *all* exchanges in which bid reductions were sufficiently profitable, regardless of whether the exchange participated in header bidding. Google's bidding behavior on its own AdX exchange is also consistent with this same principle. Before 2019, bid reductions on AdX were not sufficiently profitable to trigger Poirot. Since 2019, AdX uses a first-price auction, and DV360 computes the optimal bid shading for its advertisers bidding on AdX using the Poirot algorithm.[488]

246. *Second*, Poirot did not shade bids into exchanges that used second-price rules, regardless of whether they participated in header bidding. Typical Poirot experiments did not find any benefits to altering bids on exchanges such as Improve Digital and United, which appeared to be second-price exchanges.[489] Google's internal study found that Poirot

---

[487] Fourth Amended Complaint ¶ 400; Expert Report of J. Gans (Jun. 7, 2024), at ¶¶ 863-64 ("In addition to the conduct that I have found to be anticompetitive in itself, there were additional actions aimed at limiting the impact of Header Bidding [...] As one example, Google increased barriers to entry through Projects Poirot and Elmo.").

[488] *See* Presentation, "DV360 optimizations ENG deep dive" (Jan. 24, 2020), GOOG-DOJ-11733552, at -579 ("The concepts of Project Poirot were leveraged to let DV360 bid into the AdX first-price auction"); Declaration of N. Jayaram (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶¶ 35-36 ("Poirot was also updated following Google's shift to a Unified First Price Auction.").

[489] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -641 to -642. These were either second-price auctions, or sufficiently close to second-price auctions that Poirot did not typically detect a sufficiently large benefit to bid shading.

would increase spending on such non-Google second-price exchanges by [redacted]%.[490] This further suggests that Poirot benefited some competing exchanges.

247. *Third*, Poirot shaded bids only as much as necessary to maximize advertiser surplus and no more, as an objective to "combat" what other exchanges might predict.[491] A program designed to "combat" an exchange might shade bids below the optimum or choose not to bid on it at all, but that is not the effect that Poirot had. Moreover, as I discuss in the subsequent paragraphs, if an exchange changes its auction format and bidders adapt to that change by choosing surplus-maximizing bids, there is no reason to expect *a priori* that the exchange's revenue will fall at all.

3. Plaintiffs' and Their Experts' Theory of Harm to Publishers and Competing Exchanges Is Based on Faulty Assumptions

248. Plaintiffs claim that Poirot aimed to limit the scale of competing exchanges.[492] Plaintiffs also point to internal documents projecting that "non-second price exchanges [would] see a revenue drop in the range of [redacted]% … Overall [DV360] revenue impact is [redacted]%" as part of a proposed theory of harm to publishers and competing exchanges.[493] Professor Gans similarly argues that "Poirot [...] made it more difficult for rival exchanges and

[490] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -815 to -816 ("Assuming that we continue to spend all budgets, the following table shows the impact once the budget server adapts. […] Clean 3p DBM revenue […] [redacted]%").

[491] *See* Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -811 ("Our optimization problem can be stated as follows: For each advertiser [redacted] Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -633 ("Most advertisers and agencies using DBM, however, don't have the technology and sophistication needed to combat adversarial SSPs.").

[492] Fourth Amended Complaint ¶ 401 ("Google's main goal was depriving rival exchanges of sufficient scale.").

[493] Fourth Amended Complaint ¶ 401 ("Initial experiments regarding the effect of Poirot actually showed a negative revenue impact to DV360, but Google's main goal was depriving rival exchanges of sufficient scale engaged in header bidding to compete with Google's ad exchange…").

entrants to compete," because it "resulted in reallocating revenue from rival exchanges to Google's own exchange."[494] There are two errors in the Plaintiffs' theory of harm to competing exchanges.

249. *First*, Poirot was not designed to harm competing exchanges: in fact, non-Google exchanges running second-price auctions would expect to see *increased* spending as a result of Poirot. Plaintiffs neglect to include the other finding in the document quoted in the Complaint that "[c]lean exchanges overall [would] see [redacted]% revenue increase."[495] This could be expected as a result of the fact that Poirot helped advertisers avoid overpaying for impressions, so that the spending saved by purchasing impressions at a lower price could be spent on other impressions. Google engineers expected this would occur, writing, "we expect the budget server to react quickly to adjust impression probabilities and bring [budget-constrained advertisers] back up to spending their budgets."[496] Indeed, post-launch studies found that the projected revenue drop "didn't materialize after launching," which is consistent with advertisers reallocating cost savings to the purchase of additional impressions. After incorporating these changes, Google studies showed that Poirot increased spending on competing second-price exchanges by [redacted]%. Even if advertisers chose not to spend the Poirot cost savings in this way, Poirot still increased the average profitability of each impression acquired, resulting in a net increase in total

[494] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 865.

[495] Design Doc, "Project Poirot" (Mar. 31, 2017), GOOG-DOJ-11247631, at -631.

[496] Design Doc, "Poirot Design Doc" (Apr. 25, 2017), GOOG-DOJ-13627809, at -814. To understand how this effect could occur, imagine that, without Poirot, an advertiser had a $100 budget and won 1000 impressions, for an average cost per impression of $0.10. Optimal bidding under Poirot into an exchange using non-second-price auctions might allow the advertiser to win the same 1000 impressions for a lower average cost per impression of, say, $0.09, leaving $10 left in the advertiser's budget, which could be used to purchase an additional 111 impressions at that same average cost.

advertiser surplus.

250. *Second,* any harms to non-second-price exchanges rely on the Plaintiffs' persistent (but faulty) assumption that advertisers do not change their bids in response to changing incentives. Poirot was a free service to DV360 advertisers and, in its absence, advertisers would be incentivized to pursue their own bid optimization programs and these would most likely be less efficient than the one that DV360 designed. Such self-service is the relevant comparison to assess the effects of programs like Poirot, not the Plaintiffs' fictitious but-for world in which advertisers do not respond to incentives.

251. In the absence of programs like Poirot, advertisers that continued to report the same fixed CPMs to DV360 despite exchanges transitioning away from second-price auctions would experience large reductions in their profits from online display advertising. For example, an advertiser that used DV360 to bid into an exchange that suddenly transitioned from using a second-price auction to using a first-price auction would notice—in the absence of Poirot—that it suddenly started paying more for impressions than it did before: in particular, a price equal to its bid. If that bid was optimized for a second-price auction, so that it equaled its value for the impression, the advertiser would suddenly find itself earning *zero* advertiser surplus on the impressions it won. Clearly, that advertiser would be incentivized to respond, either by excluding that exchange or by reducing the fixed CPM it reports to DV360 to use for bidding.

252. Additionally, without a service like Poirot, each advertiser would face the complex task of identifying optimal bids on its own, which would require costly experimentation and engineering resources. Such experimentation was made more complicated by the

presence of exchanges using dirty auctions, which sought to obscure their auction rules. Even so, at around the time Poirot was introduced, other buying tools had already started to develop bid optimization programs for non-second price exchanges.[497] Many DSPs including [redacted], The Trade Desk, MediaMath, [redacted], [redacted] implemented programs similar to Poirot, indicating that bid optimization is perceived as a valuable service for advertisers.[498]

253. Poirot made bidding easier for DV360 advertisers by performing bid optimizations for them, and more efficiently. DV360 and other DSPs could adapt bids to more information than any single advertiser observes, allowing them to bid more efficiently on behalf of their advertisers than any individual advertiser could do on its own. In this way, Poirot (and similar programs used by other buying DSPs) reduced both the costs to advertisers of optimizing their bids and the likelihood of bidding errors.

---

[497] *See* Sarah Sluis, "Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price," AdExchanger (Sep. 5, 2017), https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/ ("To combat price increases, some buyers have already started shading, or reducing bid prices. But that strategy comes with its own risks: Buyers will lose out on inventory they want if they submit too low a bid and the auction turns out to work with second-price logic.").

[498] *See* Presentation, [redacted] *See also* Sarah Sluis, "Everything You Need To Know About Bid Shading," AdExchanger (Mar. 15, 2019), https://www.adexchanger.com/online-advertising/everything-you-need-to-know-about-bid-shading/ ("The Trade Desk charges a fee for Koa, its bid shading algorithm, where it pockets a percentage of how much buyers save on each impression. MediaMath also developed bid shading capabilities.").

254. To assess any harms of Poirot to publishers and non-second-price exchanges, a more revealing question is whether Poirot caused the payments to exchanges using non-second-price auction rules to fall, on average, below the payments for similar impressions on a second-price exchange. Auction theory suggests that the answer is "no." In the standard independent private values auction model (adopted by Plaintiffs' experts for their analysis[499]) the celebrated **Revenue Equivalence Theorem** implies that the average revenue for publishers and exchanges is the same with profit-maximizing bids (which Poirot was designed to achieve) in a first-price auction as in a second-price auction.[500] Poirot helps advertisers to bid optimally but would not be expected to reduce other exchanges' average prices below the same market-clearing level that is achieved by second-price auctions.

255. I present an especially simple example to illustrate the power of this *revenue equivalence* idea. Suppose that there are two bidders: Bidder 1 with a value of $1.00 CPM for each impression, and Bidder 2 with a value of $2.00. Each bidder knows its own value but not the other's value. Suppose the publisher floor price is less than $1.00.

256. In a second-price auction, both bidders behave optimally by bidding their values. The seller is paid the value of the second-highest buyer, which is $1.00. In a first-price auction, both bidders could experiment to learn how to bid optimally. Bidder 1 will never experiment by bidding higher than its value of $1.00, because any higher bid can never

[499] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...] and it is a sensible assumption to make").

[500] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 73-77. *See also* Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58-73; Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, *13*(3), 227-86; McAfee, R. P., & McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, *25*(2), 699-738.

yield a positive surplus. Bidder 2's experiments will reveal that it always wins the impression when it bids at least $1.00, so it will learn not to bid more than that. It may also experiment with bids less than $1.00 and find that those bids too often lose the auction. This process of experimentation would ultimately result in Bidder 2 learning to bid just above $1.00 to always win the auctions. The resulting per-impression publisher revenue in the first-price auction will be around $1.00, the clearing price in the second-price auction.

257. For both kinds of auctions, the advertiser with the $2 value eventually learns to bid optimally, winning most of the impressions and paying about $1. Although my example is very simple—it includes only first- and second-price auctions and makes learning easy with its assumption that values do not change from impression to impression—it illustrates a general principle: bid adjustments by optimizing bidders may fully offset the direct effects of changes to auction rules. Omitting bidders' attempts to optimize can mislead, so a full analysis must account for incentives to do so.

4. Poirot Applied Equally to AdX, as Well as to Other Exchanges

258. Plaintiffs claim that Poirot reallocated advertiser spending from dirty auctions to AdX despite it "engaging in the very same auction manipulation."[501] But Poirot applied the same [redacted]% profit threshold to AdX as it did to other exchanges and found that bidding truthfully was an optimal or near-optimal bidding strategy on AdX, and accordingly did

[501] Fourth Amended Complaint ¶ 400 ("Although DV360 was openly critical of 'greedy' rival exchanges that claimed to run a true second-price auction while actually running a 'dirty' second-price auction, Google's own exchange was engaging in the very same auction manipulation [...] Accordingly, DV360 intentionally bid less on rival exchanges and increased bids on its own ad exchange, ostensibly to avoid optimizations that were bad for advertisers, when DV360 was actually redirecting that ad spend to a marketplace that engaged in exactly the same behavior.").

not shade its AdX bids. Similarly, Poirot did not adjust bids into non-Google exchanges such as United and Improve Digital, where optimal bid shading was also predicted to increase advertiser surplus by less than [redacted]%.[502] Indeed, Poirot acted as a limit on AdX's designs, with AdX deciding to reject a version of its Reserve Price Optimization program (RPO) in order to avoid triggering "things like Poirot."[503]

259. Figure 8 displays results from "typical" Poirot v1 experiments, which show the gains from bid shading on different exchanges.[504] While DV360 could add more than [redacted]% to advertiser surplus by shading bids into Pubmatic and OpenX, shading bids into AdX and Improve Digital would add less than [redacted]%. Subsequent versions of Poirot made similar findings.[505]

[502] See typical experiment on these exchanges displayed in Figure 8.

[503] *See* "AdX Dynamic Price Meeting Notes" (Apr. 11, 2018), GOOG-AT-MDL-012701069, at -073 ("Online RPO [:] Three potential less aggressive versions [;] Waiting on gTrade team to let us know which versions are acceptable for Poirot"), -075 ("Online RPO [:] Working on less aggressive version to avoid things like Poirot").

[504] *See* Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -641.

[505] *See* Design Doc, "Poirot v2.0" (Aug. 10, 2018), GOOG-DOJ-12059682, at -683 to -684 ("The new poirot [sic] model turns out to be a no-op on AdX similar to the prod model.").

Figure 8: Bid multipliers and advertiser surplus[506]

[redacted]

260. The designs of Poirot are consistent with DV360 offering a service to maximize surplus for its advertiser-customers, rather than to preference Google's AdX exchange. As I discussed in Paragraph 226 above, the statistical threshold used in Poirot made sense as a way to avoid reducing bids on genuine second-price auctions, which would reduce advertiser surplus and might inadvertently harm exchanges using a second-price auction.

[506] Presentation, "Bidding in adversarial auctions" (Nov. 27, 2017), GOOG-DOJ-05282625, at -641.

VIII. DYNAMIC ALLOCATION: USING AUCTIONS TO INCREASE PUBLISHER REVENUES ON REMNANT IMPRESSIONS

A. Overview

261. Dynamic Allocation (DA) was an auction design introduced by DoubleClick in 2007 to improve publishers' sales of remnant impressions.[507] At the time of its launch, DA differed from the dominant payment models of ad networks and the waterfall—a common method used by publishers to allocate remnant impressions—in its use of a second-price auction to determine payments to publishers.[508] Under DA, a publisher could configure DFP to sell a non-guaranteed impression to a bidder on AdX only when it would pay more than the publisher's largest expected payment from any other remnant demand source.[509]

262. As a consequence, on every impression, DA could only increase a publisher's expected revenues compared to a *status quo ante* without bids from AdX. To quantify the effects of DA for publishers and evaluate these effects relative to a wider range of counterfactuals, I conducted a simulation study using GAM and Google Ads auction data from January of 2024. The simulations suggest that average publisher revenues using DA are substantially higher than those obtained using a waterfall, both in the case that AdX is not included in

[507] *See* Presentation, "2008 Strategic Planning DoubleClick Advertising Exchange" (Jul. 26, 2008), GOOG-TEX-00458239, at -247; Presentation, "Ad Exchange Dynamic Allocation" (Sep. 5, 2013), GOOG-TEX-00054839, at -843.

[508] *See* White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 to -414.

[509] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -414 ("If [AdX] can provide the publisher with a net CPM value higher than they would have gotten from delivering their directly booked, non-guaranteed ad, [AdX] will deliver an ad.").

the counterfactual waterfall (a [redacted]% increase) and in the case that AdX would otherwise have been included in the counterfactual waterfall (a [redacted]% increase).[510] In my simulations, DA also expanded output by reducing the number of unsold impressions.

263. Plaintiffs' and their experts' arguments about DA have significant flaws:

a. Plaintiffs and their experts allege that DA steered transaction volume towards AdX.[511] But DA only accomplished that via competition on the merits, delivering what publishers wanted: substantially higher revenues.

b. Plaintiffs' allegations lack historical context, comparing DA to an ahistorical counterfactual in which other exchanges could make real-time bids into DFP.[512] DA improved publisher revenues and allocated impressions more efficiently than the waterfall model that preceded it, which did not make any use of real-time bids.[513] The technology to accept real-time bids from AdX bidders was introduced in 2009 (after DoubleClick's acquisition by Google[514]) with the launch of "AdX

[510] I describe these simulations in detail in Section VIII.E and the technical notes in Section XV.B.

[511] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 548 ("In its implementation of DA after the acquisition of DoubleClick, Google made and maintained critical choices with the intention of steering inventory to its AdX exchange compared to other intermediaries, without providing benefits to publishers."). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 120 ("In my opinion, Dynamic Allocation led to higher win rate and higher revenue for AdX as well as lower win rate and lower revenue for non-Google exchanges."); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 123; Fourth Amended Complaint ¶ 281("Dynamic Allocation was exclusionary and successfully foreclosed competition in the exchange and buying tool markets").

[512] *See, e.g.*, Fourth Amended Complaint ¶ 279 ("Internal Google documents reveal Google's knowledge of its own misrepresentations, stating that the optimal publisher set up in display advertising includes 'real-time bidding across exchanges,' which is 'at scale, at the best possible price, with zero waste.'").

[513] *See* White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("With indirect sales, the CPM is usually fixed, but the number of impressions delivered is not.").

[514] Google, Google Closes Acquisition of DoubleClick, News from Google (Mar. 11, 2008), http://googlepress.blogspot.com/2008/03/google-closes-acquisition-of_11.html ("Google Closes Acquisition of DoubleClick").

2.0," improving efficiency and publisher revenues further.[515] Header bidding, a technology to compare real-time bids from multiple exchanges (discussed further in Section X) did not gain popularity around 2014,[516] years after Dynamic Allocation launched. Designing an "auction of auctions" to compare the highest real-time bids from multiple exchanges natively on Google Ad Manager presented additional challenges, including changing the auction design and helping bidders adapt to that, and Google solved those problems with its rollout of a Unified First Price Auction in 2019 (which I discuss further in Section XIII). Plaintiffs ignore this evolution of the industry over more than a decade and make comparisons of DA technology created in 2007 to outcomes that were possible in 2019.

c. Plaintiffs' and their experts' conclusions are based on analyses that fail to sufficiently account for publishers' incentives to optimize floor prices: when these incentives are properly accounted for, those conclusions change.

B. Online Display Advertising Configurations Prior to DA

264. In the early years of online display advertising, most revenue from online display advertising inventory was generated via negotiated **guaranteed contracts** that fixed a

---

[515] *See* Email from S. Woods to S. Feldman, "Re: [Adsense-eng-wat] [Adsense-eng] Re: [Ads-engdirs] Doubleclick Ad Exchange 2.0 - Launched!" (Sep. 19, 2009), GOOG-AT-MDL-010836318, at -318 ("The team has done a great job […] to also go beyond in some very important areas, e.g. […] Real Time Bidding"); White Paper, "DoubleClick Ad Exchange Impact" (Q4 2010), GOOG-DOJ-13247322, at -322 ("The results of our research showed that when DoubleClick Ad Exchange wins the auction, publishers generate [redacted] on average, net of revenue sharing and ad serving fees, compared with fixed upfront sales of non-guaranteed display advertising. […] The Ad Exchange fill rate (percentage of offered impressions resulting in a matched transaction, including those with minimum CPMs or restrictions) [redacted]

[516] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("Around 2014, web publishers began to adopt Header Bidding."). *See also* AdPushup, "Header Bidding" (2023), https://www.adpushup.com/header-bidding-guide/ ("Header bidding made it to ad tech somewhere around 2014. And only after one year, in 2015, the technique went viral.").

price and quantity of certain types of impressions.[517] Publishers monetized their remnant impressions (impressions not allocated to guaranteed contracts) using **indirect** sales channels, which in the early years of online display advertising consisted primarily of ad networks.[518] Ad networks would purchase remnant inventory from publishers and package it for sale to advertisers. Publishers would sell advertising space to ad networks on a per-impression basis, with no obligation to fill a minimum number of impressions.[519] Many ad networks paid publishers a fixed price per impression or a fixed proportion of the revenue it earned from advertisers,[520] but did not give publishers control over the ads that would be placed on their websites or the prices they would receive for impressions.[521]

---

[517] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("A typical large publisher generates upward of [redacted]% of its online advertising revenue from guaranteed ad sales.").

[518] *See* White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("To tap additional demand from advertisers, brand-conscious publishers often sell through indirect channels, including ad networks, exchanges, and other technology providers").

[519] *See* White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Publishers usually sell their remaining ad space on a non-guaranteed (or 'pre-emptible') basis through their direct sales channel, as well as through their indirect sales channel, which may comprise a handful of ad network partners. […] With indirect sales, the CPM is usually fixed, but the number of impressions delivered is not."); "AFC Partnerships" (Jan. 7, 2008), GOOG-DOJ-03516570, at -575 ("Today, most large publishers […] rely on ad networks like AFC to fill remnant inventory.").

[520] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Ad networks provide revenue to publishers willing to commit, in advance, a large number of impressions at a fixed CPM value.").

[521] *See* David Kaplan, "On Ad Networks: Pork Bellies, Diamonds, Or The New Direct Marketing?," Forbes (Apr. 8, 2008), https://www.forbes.com/2008/04/08/online-ad-networks-tech-cx_pco_0408paidcontent.html?sh=7414ef02cb8e ("All ad networks are not created equal: If all sides can agree on one thing, it's the need for greater clarity to what's being sold and where it's being placed. […] 'In a lot of cases [in terms of ad nets' handling of remnant, or unsold ad inventory], the buyer doesn't really know what they're getting. And the seller doesn't have any control over price.'"). *See also* "Ad Networks," AffiliateSeeking.com (captured on Jan. 16, 2008) https://web.archive.org/web/20080116101025/https://www.affiliateseeking.com/list/23000001/1.html (listing some examples of ad network pricing options).

265. At times, an ad network might not have a relevant ad to fill a publisher's impression. Ad networks developed the capability to **passback** an unwanted impression, allowing the publisher to offer it to other demand sources.[522] This capability led to the **waterfall**, in which a publisher specifies a list of demand sources to be sequentially offered the opportunity to fill an impression.[523] Under a waterfall, the first demand source was offered the opportunity to fill an impression. If that demand source declined to fill the impression, the request was passed back to the next demand source on the publisher's list, with the process repeating until the impression was sold or the publisher's list was exhausted (or a timeout limit was reached), leaving the impression unsold. The waterfall was a highly configurable process, with publishers free to set the order of consideration and floor prices for each demand source however they wished. Because publishers often did not know whether an ad network would have an eligible ad for the given impression, or what the ad network would pay if it did have an eligible ad, a common approach was to prioritize ad networks in the waterfall with the highest historical average payouts per impression.[524]

[522] Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -503.

[523] Maciej Zawadziński and Mike Sweeney, "What is Waterfalling and How Does it Work?," Clearcode Blog (Aug. 20, 2021Sep.), https://clearcode.cc/blog/what-is-waterfalling/ ("Waterfalling gets its name from the waterfall-like process for selling inventory—i.e. the demand sources are initiated one at a time, one after another. […] the publisher first tries to sell its inventory via direct sales, as these generally offer the highest cost-per mille (CPM). If it is unable to do so, the publisher will then pass the impression down the waterfall to various ad networks until it is sold."); Presentation, "Ad Manager Ecosystem 101" (Jun. 2019), GOOG-DOJ-AT-02199478, at -503.

[524] *See* Maciej Zawadziński and Mike Sweeney, "What is Waterfalling and How Does it Work?," Clearcode Blog (Aug. 20, 2021Sep.), https://clearcode.cc/blog/what-is-waterfalling/ ("Publishers, however, started running into problems when their chosen ad network wasn't able to sell all of their remnant inventory, which meant their available ad spaces were being left unfilled, leading to missed revenue opportunities. In an effort to increase fill rates and capitalize on revenue opportunities, a process known as waterfalling started to emerge."); White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Some publishers manage yield across this ad space by manually prioritizing which ad networks access it in order of their relative average CPM payout.").

266. In a waterfall, if a publisher does not know which demand sources are interested in purchasing any particular impression but knows exactly the prices each interested source would pay, then it can maximize its revenue by ordering the waterfall list from the highest price offer to the lowest. But when the prices from different demand sources are not known, for example because they are set by separate auctions within each ad network, then no waterfall process can guarantee selling to the highest bidder: the publisher cannot know whether another buyer lower on its list would have bid more. With uncertainty about price offers, the publisher is incentivized to reject some bids from demand sources near the top of the waterfall, even if the same bids would be accepted from sources near the bottom of the waterfall. The publisher can implement that policy by setting *higher* minimum prices nearer the top of the waterfall.

267. Even so, by offering each impression *sequentially* to ad networks, the waterfall procedure could leave value on the table. Under the waterfall, an impression might be assigned to an advertiser on Ad Network A with a lower value for the impression than an advertiser on Ad Network B, merely because Ad Network A had a higher priority in the publisher's waterfall. That was in large part because, at that time, the prevailing technology did not permit an efficient "real-time" auction among different demand sources. Publisher revenue and advertiser surplus could *both* increase by reallocating the impression to the advertiser on Ad Network B at some higher price than Ad Network A paid, if the technological limitations that existed at the time could be overcome and standards for interoperability could be established.

C. DA Introduced Auctions for Impressions, Improving Platform Thickness and Increasing Publisher Revenues

268. **Dynamic Allocation**, introduced in 2007,[525] improved the allocation by replacing the *sequential* logic of the waterfall with the *simultaneous* comparison of bids from advertisers, ad agencies and ad networks participating in a real-time auction on AdX. In the version of AdX that existed when DA was introduced, participating demand sources would enter bids in advance of auctions, with each bid containing targeting criteria that told AdX the types of impressions that the demand source was interested in purchasing at the specified bid amount.[526] In DFP, publishers would codify information about their non-AdX sources of remnant demand using non-guaranteed line items, which also contained targeting criteria and a **value CPM** (also known as a **static bid**), which DFP used to represent that source of demand in the DA process.[527]

269. DA used a two-step procedure to allocate remnant impressions. First, it would identify the eligible non-guaranteed line item with the highest value CPM: Google engineers called this static bid the **DFP booked price**.[528] Then, AdX would run a second-price

---

[525] Presentation, "2008 Strategic Planning DoubleClick Advertising Exchange" (Jul. 26, 2008), GOOG-TEX-00458239, at -247 ("Q3'07 Dynamic Allocation.").

[526] DoubleClick, "Ad Selection Specifications for Ad Server Version 14.1" (Mar. 27, 2007), GOOG-AT-MDL-007374059, at -136 ("Buyers can bid on ad slot inventory by indicating their preferred targeting elements and specifying an associated CPM based bid.").

[527] In the original version of AdX, prior to the acquisition by Google, the terminology was a little different: line items were just 'ads' and value CPMs were just 'bids,' but to avoid confusion with other places in this report, I will adopt Google's terminology. *See* DoubleClick, "Ad Selection Specifications for Ad Server Version 14.1" (Mar. 27, 2007), GOOG-AT-MDL-007374059, at -136; DoubleClick for Publishers, "Terminology differences with DART," DoubleClick for Publishers Help (captured on Feb. 11, 2012), https://web.archive.org/web/20120211164005/http://support.google.com/dfp_premium/bin/answer.py?hl=en&answer=158833.

[528] Presentation, "Ad Exchange Dynamic Allocation" (Sep. 5, 2013), GOOG-TEX-00054839, at -843 to -854; Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶¶ 10-12.

auction using the bids it received from advertisers on AdX with a floor price at least as high as the DFP booked price.[529] The AdX bidder with the highest bid would be allocated the impression, as long as that bid was above the floor price; otherwise, the impression was allocated to the demand source associated with the best non-guaranteed line item. If that demand source had no ad to serve, it might pass the impression back to other demand sources in a waterfall. DA was backward compatible: a publisher could adopt DA without major changes to its existing relationships with ad networks.

270. By using a second-price auction to allocate impressions, DA eliminated the possibility of inefficient allocation among the bidders in the auction and ensured that the publisher received a price for the impression that no other AdX bidder was willing to beat. But DA had another important feature in its design in recognition of the fact that publishers had other potential demand sources for their remnant impressions: a floor price in the form of a value CPM set by the publisher.[530] A publisher would be incentivized to set those value CPMs so that the effective floor price for each impression was at least as large as the expected revenue from any non-AdX demand source, ensuring that a bidder on AdX could win only if it paid at least that amount. As long as the publisher did so, DA could increase the expected revenue to the publisher from each impression it sold, but could never reduce it.

[529] If the publisher had otherwise set a floor price for the impression that was higher than the DFP booked price, that floor price would apply. *See* "Adx Queries by Pricing Rule" (Sep. 1, 2016), GOOG-DOJ-13470118, at -118 ("For every query in adx, the price can be determined by the following: [...] PUBLISHER_RESERVE: the reserve set by the publishers[.]").

[530] As Professor Weinberg notes, in addition to using value CPMs, publishers set an explicit floor for AdX; when describing how DA works he writes: "Every lower priority static line item, including AdX, has [...] a price floor [...] DFP calls AdX with reserve price equal to the maximum of [the reserve derived from value CPMs] and AdX's price floor." Expert Report of M. Weinberg (Jun. 7, 2024) at ¶ 113. It is not important for my analysis which mechanism publishers use to set the binding floor prices.

271. A publisher using DA would experience three immediate benefits. *First*, if the publisher was not already calling AdX, then DA allowed it to gain access to a new source of advertising demand. New advertising demand meant new opportunities for publishers to sell remnant impressions, and fewer unsold impressions. *Second*, if AdX had two bidders willing to beat the floor price for DA, then the auction price would exceed the floor, further increasing seller revenues. *Third*, DA gave the publisher a more efficient allocation method to sell display advertising inventory, guaranteed to increase its expected revenues from the sale of remnant impressions compared to using the waterfall with non-Google demand sources alone. By choosing a floor price for each line item that is higher than the expected payment from alternative demand sources (*i.e.*, any demand sources listed in the waterfall that might be triggered if the impression is not sold on AdX), the publisher would earn a higher expected revenue from each impression sold on AdX. This property made DA a *risk-free revenue improvement* for publishers integrating AdX as a new source of demand. An analysis of the first version of DA found that the combined effects of these three benefits of DA increased the revenues that publishers earned per impression sold on AdX by [redacted]%.[531]

272. DA also benefited the advertisers and ad agencies participating in the auctions on AdX: in addition to having access to additional inventory from publishers, they could always win an impression by bidding enough for it, rather than relying on the comparatively non-transparent processes used by ad networks to buy impressions on their behalf.[532]

[531] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -415 ("The results of our research demonstrated that the combined effects of auction pressure and Dynamic Allocation in DoubleClick Ad Exchange resulted in an average CPM lift of [redacted]% compared with fixed, upfront, pre-negotiated sales of non-guaranteed inventory.").

[532] *See* David Kaplan, "On Ad Networks: Pork Bellies, Diamonds, Or The New Direct Marketing?," Forbes (Apr. 8, 2008),

Moreover, the second-price auction ensured that, on each impression won by AdX bidders, the winning bidder paid *only* the amount needed to beat other bids on AdX and the floor price determined by the publisher, and not more, which as I explained before, made bidding simpler for advertisers (see Section III.C.3.a).

D. Real-Time Bidding in AdX 2.0 Further Increased Benefits of DA for Publishers and Advertisers on AdX

273. In September 2009, after its acquisition of DoubleClick, Google redesigned the AdX exchange to further increase the benefits of online display advertising auctions for publishers and advertisers. The redesigned **AdX 2.0** incorporated **real-time bidding** from bidders on AdX and Google's buy-side products (originally just AdWords, but later also DV360).[533] Under real-time bidding, AdX would calculate real-time bids from Google's buy-side products and send a **bid request** containing information about the impression to non-Google bidders on AdX (later called **Authorized Buyers**). After receiving a bid request, bidders would use real-time information (such as ad campaign information and cookies) to determine their bids for the impression, which they would send to AdX.[534]

https://www.forbes.com/2008/04/08/online-ad-networks-tech-cx_pco_0408paidcontent.html?sh=7414ef02cb8e ("All ad networks are not created equal: If all sides can agree on one thing, it's the need for greater clarity to what's being sold and where it's being placed. […] 'Both buyer and seller require transparency. In a lot of cases [in terms of ad nets' handling of remnant, or unsold ad inventory], the buyer doesn't really know what they're getting. And the seller doesn't have any control over price.'").

533 Email from [redacted] to [redacted], "Re: [Adsense-eng-wat] [Adsense-eng] Re: [Ads-engdirs] Doubleclick Ad Exchange 2.0 - Launched!" (Sep. 18, 2009), GOOG-AT-MDL-010836318, at -318 ("I am extremely thrilled to see AdX 2.0 launch. The team has done a great job not only getting to parity with AdX 1.0 but to also go beyond in some very important areas, e.g. API support, Real Time Bidding, and of course integration with Adsense and Adwords.").

534 Maciej Zawadziński, "How Does Real-Time Bidding (RTB) Work?," Clearcode Blog (Jul. 2, 2021), https://clearcode.cc/blog/real-time-bidding/ ("The supply-side platform analyzes the information about the user (location, web history, and, if available, age, gender and any other user information) and then sends this information to the ad exchange. Once the ad exchange receives this information, it connects to the demand-side platforms and relays information about the user. The ad exchange starts an auction, and the DSPs then bid on the impression based on what that particular impression is worth to them -determined by predefined parameters set by the advertisers.").

AdX would then run an auction using those bids and return any winning ad to DFP, which served code for the winning ad to the publisher's website. This entire process—from the arrival of the impression to the collection and processing of bids and the presentation of any winning ad—would be completed in the blink of an eye,[535] avoiding slow page load times for end users. Initial reactions from industry players to the launch of AdX 2.0 heralded it as a "watershed moment in the progression towards truly dynamic, demand-driven advertising transactions" and a "positive event for anyone in the exchange space."[536]

274. Real-time bidding improved publisher revenues and the matching of impressions to advertisers. Static bids could leave value on the table for both publishers and advertisers: an advertiser might be willing to pay much more for an impression about which it had accurate real-time information (for example, whether the user had recently visited its website) than it might offer using static bids determined without real-time information. Real-time bidding also allowed a bidder to develop its own methods for determining bids, using any additional information the bidder might have about the impression that it was unable to incorporate into any static bidding system. For this reason, real-time bidding offered potential benefits to both publishers and advertisers.

[535] Google, "How Authorized Buyers Work With Google Ad Manager," Google Ad Manager Resources (accessed Sep. 27, 2023), https://admanager.google.com/home/resources/how_authorized_buyers_work_with_google/ ("This all happens within around 100 milliseconds.").

[536] Email from [redacted] to [redacted], "FW: comments from industry players on AdX 2.0 on AdExchanger.com this evening" (Sep. 22, 2009), GOOG-AT-MDL-B-003180112, at -113 ("I think everyone is excited to see how [Google] can apply their expertise to the next generation of biddable display."), -113 ("The rollout of Google's updated DoubleClick Exchange offering - and the proliferation of other similar real time bidding platforms - marks a watershed moment in the progression towards truly dynamic, demand-driven advertising transactions."), -112 ("I think the launch of AdX 2.0 is an example of growing the pie as opposed to stealing share from a competitor because it's going to bring lots of new sellers (AdSense and DART for Publisher sites) and buyers (anyone who uses AdWords) into the market. As those sellers and buyers get comfortable with the exchange model I think they'll begin trading on the other platforms as well, so I think this is a positive event for anyone in the exchange space.")

275. To illustrate the potential benefits of real-time bidding over payments based on static bids, suppose that a publisher's best offer from an ad network, say Network X, was a fixed $1 CPM, while an AdX bidder was willing to pay 70¢ three-quarters of the time and $1.50 otherwise. Under DA, the publisher would be incentivized to choose a floor price for AdX of at least $1.00, its best offer from a competing ad network. If the AdX bidder was not able to make real-time bids for the impression, its optimal static bid would be equal to its average value of 90¢.[537] That bid would never win. If instead the AdX bidder can make real-time bids, it could win the impression when its value was $1.50, while the ad network would win the other three-quarters of the time. Whether the publisher, the AdX bidder, or both parties benefit from real-time bidding in this example depends on the floor price chosen by the publisher. With any floor price between $1.00 and $1.50, both the publisher and the AdX bidder benefit.[538]

276. Contrary to Professor Weinberg's conclusion that DA "led to [...] a lower win rate and lower revenue for non-Google exchanges,"[539] some of the benefits of real-time bidding can accrue to non-Google demand sources. To show that, consider the following modification of the previous example. Suppose that Network X continues to offer a fixed $1 CPM payment but the highest AdX bidder instead has a value of $1.50 three-quarters of the time and 70¢ otherwise, reversing the previous probabilities. In the absence of real-time bidding, its optimal static bid would be its average value of $1.30[540] for each

[537] Its average value is 3/4 times 70¢ plus 1/4 times $1.50, which equals 90¢.

[538] With a floor price of $1.00, the publisher would make no additional revenue while the AdX bidder earns surplus from the ads it wins when its value is $1.50. With a floor price of $1.50, publisher revenues increase, while the AdX bidder pays its value, leading to no additional profits.

[539] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 101.

[540] In this case, its average value is 3/4 times $1.50 plus 1/4 times 70¢, which equals $1.30.

impression, which would always win the auction over the static bid of $1 from Network X. After real-time bidding is enabled, the AdX bidder would only win impressions when it had the higher value of $1.50 (three quarters of the time). In that case, Network X—the network without real-time bidding—would benefit from the transition to real-time bidding, because it would win the impression in the one quarter of cases when the AdX bidder has a value of 70¢. The publisher can also benefit from the transition to real-time bidding in this example by increasing its floor price in the AdX auction above $1.00, extracting more revenues from the AdX bidder when it has a high value for the impression.

277. The preceding examples understate the benefits of DA to publishers since, for simplicity, I have assumed there is only one bidder on AdX. In more realistic examples with multiple AdX bidders, the benefits to publishers would be higher because the AdX winner's price would be the maximum of the floor price and the second-highest bid.

278. To realize the benefits of real-time bidding when many ad networks offered fixed-price payments or revenue sharing to publishers, ad servers needed to adapt their allocation methods to integrate both approaches—a process that was widely understood in the industry to be challenging. OpenX cofounder Jason Fairchild summarized the challenge of unifying live and static demand, saying, "If you think about it, they're two fundamentally different marketplaces. To combine them, you have to rethink even your auction mechanism […] Everything is radically different."[541] DA with AdX 2.0, launched

[541] Josh Ong, "Adtech firm OpenX unveils an industry-changing fusion of real-time bidding and ad networks," The Next Web News (Jun. 9, 2014), https://thenextweb.com/news/adtech-firm-openx-unveils-industry-changing-fusion-real-time-bidding-ad-networks.

in 2009, was Google's answer to this challenge, and other competing exchanges introduced real-time bidding at around the same time.[542]

279. DA with real-time bidding on AdX was an important innovation offering large benefits for publishers. A Google experiment from 2010 found that DA with real-time bids offered publishers an average revenue increase on impressions sold on AdX of [redacted]% compared to pre-negotiated sales of remnant impressions.[543] On average, the experiment found that DA increased *total* publisher revenues over all remnant impressions by [redacted]%.[544] In 2013, Google estimated that DA increased publisher revenues by [redacted]% on each impression it won over an alternative remnant demand source, leading to an overall increase in publisher revenues of [redacted]%.[545] Even so, publishers could have, if they wished, disabled real-time bidding from AdX to restore their pre-DA configurations.[546]

[542] *See* Mike Nolet, "RTB Part II: Supply supply supply!," Mike On Ads (Sep. 19, 2009). http://www.mikeonads.com/2009/09/19/rtb-part-ii-supply-supply-supply/ ("Over the past few months pretty much any aggregator of supply has launched, announced or started work on some sort of RTB capability. All major exchanges — Yahoo's Right Media, Microsoft's AdECN and Google's AdEx have RTB integrations in the works. Of the pub aggregators, AdMeld & PubMatic are live and Rubicon is actively working on a solution. As mentioned, FAN has been live with Myspace inventory for a while and there are a number of other parties, such as ContextWeb, AdBrite and OpenX, entering the space.").

[543] White Paper, "DoubleClick Ad Exchange Impact" (Q4 2010), GOOG-DOJ-13247322, at -322 ("The results of our research showed that when DoubleClick Ad Exchange wins the auction, publishers generate [redacted], on average, net of revenue sharing and ad serving fees, compared with fixed upfront sales of non-guaranteed display advertising.").

[544] White Paper, "DoubleClick Ad Exchange Impact" (Q4 2010), GOOG-DOJ-13247322, at -322 ("Across all pre-emptible inventory, including those instances when the Ad Exchange did not win the auction, the revenue lift for publishers averaged [redacted]

[545] Presentation, "Ad Exchange Dynamic Allocation" (Sep. 5, 2013), GOOG-TEX-00054839, at -844 ("In these instances when the Ad Exchange won out over the alternatives, the revenue it achieved for that inventory was on average [redacted]% higher than it would have been if the Ad Exchange not been used. This translates to an average overall revenue lift of [redacted]%.").

[546] *See, e.g.*, Summer Livestream Series (Sep. 2019), GOOG-AT-MDL-B-004582905, at -3162 ("You can still exclude specific ad units from dynamic allocation if desired"); Draft Help Center Doc, "DFP and dynamic allocation" (Nov. 16, 2013), GOOG-DOJ-15416614, at -614 ("Dynamic allocation [...] allows Ad Exchange to compete in real time with line items booked in DFP. [...] Publishers configure settings in DFP and Ad Exchange in order to control which inventory is eligible to compete, and how."). Even though I am not aware of earlier references to disabling DA on DFP, it would always be possible for publishers to effectively disable DA—while still using DFP

E. My Simulation Analysis Concludes that DA Improved Outcomes for Publishers and Advertisers

280. In order to investigate the effects of DA on publisher revenues and match rates, I simulated the effects of DA on the sale of over 20 billion impressions. I informed my simulations using real-world data collected from auctions on GAM in the month of January 2024.[547, 548] I compared the outcomes from DA to those obtained under a waterfall, which was the dominant model of remnant ad allocation before the introduction of DA.

281. Simulations complement the theoretical analysis of DA presented in previous sections in two ways. *First*, simulations allow me to *quantify* the possible effects of DA, so that I can assess not only *whether* DA improved publisher revenues and ad fill rates but also *how much* it could have improved those outcomes. *Second*, simulations allow me to judge the likely qualitative effects of DA on publisher revenues and ad fill rates as compared to a wider range of counterfactuals. For example, the theoretical effects of DA I discuss in Sections VIII.C and VIII.D above compare DA to a counterfactual waterfall in which AdX does not participate. Since DA was launched around the same time as AdX, this comparison is likely to have been relevant at that time.[549] In addition, to assess Professor

---

to allocate ads to other ad networks—by setting a very high value CPM on the highest line item for DFP to beat (much higher than an AdX bidder would bid).

[547] *See* Google Ads Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000000000 to -000258388; Google Ad Manager Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000276098 to -001116097.

[548] The number of simulated impressions is calculated in [redacted] of my supporting materials, with the number saved in [redacted]. Details on each step of the simulation procedure can be found in [redacted].

[549] The DFP ad server first included DA in July 2007, shortly after DoubleClick launched its ad exchange, AdX, in North America. *See* Presentation, "2008 Strategic Planning DoubleClick Advertising Exchange" (Jul. 26, 2008), GOOG-TEX-00458239, at -246 to -247.

Gans' allegations that "[i]n its implementation of DA [...] Google made and maintained critical choices [...] without providing benefits to publishers," and Plaintiffs' allegations that DA "ultimately reduced publishers' yield,"[550] I have used simulations to compare publisher revenues under DA to a different baseline in which all demand sources, including AdX, compete on the same basis using the waterfall.[551]

282. There are several challenges associated with simulating the waterfall and second-price auctions using DA for the time when DA was introduced.

283. The first challenge is to obtain data that is informative about advertisers' values at that time. To analyze the effect of introducing DA, it would be best to have bid or valuation data from 2007, but I am not aware of any such data. To assess the effects of DA, I instead estimated bidders' values for impressions using Google Ads data and GAM auction data from January 2024. The GAM dataset contains bids for around 60 billion US impressions.[552] My analysis used a subset of this data, with my selection criteria described in the Technical Notes in Section XIII.B.2.[553] Because this dataset is from

[550] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 548 ("In its implementation of DA after the acquisition of DoubleClick, Google made and maintained critical choices with the intention of steering inventory to its AdX exchange compared to other intermediaries, without providing benefits to publishers."). *See also* Fourth Amended Complaint ¶ 274 ("Dynamic Allocation ultimately reduced publishers' yield[.]").

[551] As I discuss in Paragraph 25 and Section XIII, an alternative baseline in which all demand sources compete in a unified auction was not technologically feasible at the time of DA's launch and required additional changes to AdX's auction design. For this reason, I focus on the waterfall. Simulations based on the Open Bidding auction design would have led to even higher publisher revenues, for the reasons discussed in Section XIII.

[552] The number of impressions is calculated in code/misc_queries.py of my supporting materials, with the number saved in code/logs/misc_queries.txt.

[553] As I discuss in Section XV.B.2, I analyze the groups of auctions for which the number of observations is large enough to reliably conduct simulations. The resulting subset of data consists of auctions across [redacted] inventory units and [redacted] publishers, comprising approximately [redacted]% of relevant US publisher real-time bidding revenue in the January 2024 GAM data sample. To evaluate outcomes across a range of publisher inventory units, a separate simulation was conducted for each group of auctions. These statistics are calculated in code/misc_queries.py and saved in code/logs/misc_queries.txt. For additional details, see Section XV.B.2.

2024, my simulations could only compare the expected effect of introducing DA to a counterfactual scenario in which *contemporary* bidders instead participated in the waterfall. This may be different from what I would find if I could use data from the relevant earlier period of time. However, the heterogeneity among publishers and inventory units ensure that the data encompasses a wide variety of situations varying significantly in the relative bid strengths of various bidders and participation rates of the demand sources (including AdX).[554] Therefore, the dynamics that would have been at play in some scenarios in the relevant time period of the complaint are likely captured in at least some subset of the 2024 data.

284. The second challenge is to estimate the values advertisers used to inform their decisions in the waterfall. For Google Ads bidders, this estimation was relatively straightforward because the Google Ads dataset contains estimates of bidder values (but the data still needed to be cleaned, filtered, and matched to other auction data, as I discuss in the Technical Notes in Section XIII.B.3). For non-Google Ads bidders, however, the problem was harder because the available data contains only their bids in AdX's first-price auctions, rather than the bidders' actual values. As I have emphasized elsewhere in this report, profit-maximizing bids in first-price auctions are always strictly *lower* than advertisers' values.

[554] For example, for several publishers I simulate, Google Ads makes up a relatively small portion of the revenue coming from the demand sources I simulate: at least [redacted]% of publishers have no more than [redacted]% of their revenue coming from Google Ads, and about [redacted]% of publishers have no more than [redacted]% coming from Google Ads. These results were computed using code/misc_queries.py in my supporting materials, with the numerical values logged in code/logs/misc_queries.txt.

285. I obtained my estimates for non-Google Ads bidders' values using a state-of-the-art empirical economics technique called **bid inversion**.[555] Any method to estimate advertisers' values from the bid data must incorporate one or more assumptions about how bids and values are related. Bid inversion is based on the assumptions that (i) bidders choose the bids that maximize their expected surplus from each auction and (ii) their expectations about the distributions of others' bids correspond closely to the actual empirical distributions. Bids from Google's own demand sources, Google Ads and DV360, are generated in this way, and I would expect that other buying tools would most often choose bids in such a way.[556] An advantage of bid inversion over some other techniques used to estimate bidder values is that it is "non-parametric," which means that it does not restrict the shape of the underlying distributions of bids or values.[557] Bid inversion requires data with a sufficient density of bids in the estimation neighborhood, and that requires that I exclude intervals near the very highest bids.[558] For this study, I do not make point estimates for the highest 2% of values and restrict publishers to set their floor prices so that each demand source fills an impression at least 2% of the time that it is called to bid. I also varied the 2% threshold to 1% and 5%, to check whether other

[555] *See* Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, *68*(3), 525-574. Perrigne, I., & Vuong, Q. (2019). Econometrics of auctions and nonlinear pricing. *Annual Review of Economics*, *11*, 27-54.

[556] *See, e.g.*, Deposition of [redacted] at 369:18-21 (Sep. 17, 2021), GOOG-AT-MDL-007173084, at -453 ("A: We predict the distribution of highest other bid and then we use that information to optimize for advertiser surplus.").

[557] *See* Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, *68*(3), 525-574. The empirical data is a collection of discrete bids. Bid inversion assumes that the bidder fills the gaps between bids to create a smooth distribution of bids-to-beat before computing its optimal bid. I discuss this "bid smoothing" further in Section XV.B.4.b.

[558] This necessity is also highlighted by the cited academic studies, which do not apply bid inversion to extreme bids. *See* Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, *68*(3), 525-574. Perrigne, I., & Vuong, Q. (2019). Econometrics of auctions and nonlinear pricing. Annual Review of Economics, *11*, 27-54.

reasonable thresholds, which could lead to different floor prices, might lead to significantly different effects, and I found that my conclusions about the benefits of DA to be substantially the same. Additional technical details about my empirical approach are described in the Technical Notes in Section XIII.B.

286. The third challenge is to specify which bidders would participate using AdX and which would use other exchanges, and how those participation decisions might have been affected by the emergence of DA. I have seen no data to guide this choice, so for simplicity, I assume that, at the time DA was introduced, the only demand on AdX came from Google Ads bidders. In reality, AdX aggregated bids from many demand sources besides Google Ads, so this assumption underestimates the revenue earned by publishers from AdX in each simulation. I also assume that, to limit latency, a publisher's waterfall could use at most three participating non-Google demand sources.

287. While the first three challenges focus on valuation data, bidders, and their bids, a fourth challenge focuses on the behavior of publishers.[559] The revenues earned by publishers under DA and the waterfall depend also on publishers' floor prices and on the order in which demand sources are called. My simulations compare the waterfall's revenue to the revenue of a separately configured waterfall with DA turned on. To ensure that my estimates of DA's benefits are as conservative as possible, I identify the strongest baseline, which is when the initial waterfall is configured to maximize publisher revenue. Identifying those revenue-maximizing waterfall configurations requires finding the optimal ordering and optimal floor prices for the waterfall separately for each publisher,

[559] Determining the optimal behavior of *bidders* in the simulated AdX auction is not challenging: because each auction is conducted in a *second-price* format, the optimal strategy of each bidder in each auction is to bid its true value.

which is a challenging numerical computation. When I simulate DA, I assume (conservatively) that publishers did not optimize value CPMs exactly and instead selected the floor AdX would face by experimenting with a small set of heuristics. I describe the numerical procedures I used for this in the Technical Notes in Section XIII.B.6.

288. I measure the effects of DA under two different counterfactuals:

a. *Counterfactual 1:* In the first counterfactual, I compare calling AdX with DA to a baseline in which the waterfall would *only* call non-Google demand sources. In this counterfactual, enabling DA brings AdX as a new source of demand for publishers. In the waterfall, I include three inverted non-Google demand sources. I compare outcomes under this baseline waterfall to those that would arise if AdX were called with DA and, if AdX did not win the impression, it would be allocated through a waterfall containing just the two best non-Google demand sources.[560]

b. *Counterfactual 2:* In the second counterfactual, I compare calling AdX with DA to a baseline in which the waterfall includes both AdX (not using DA) and non-Google demand sources. In this baseline waterfall, AdX and three non-Google demand sources all participate using publisher-optimal posted prices. I compare this baseline to the outcome of calling AdX with DA, followed by a waterfall with the same three non-Google demand sources. This comparison isolates the effect of DA's introduction of auction-based pricing.

[560] I reduce the number of non-Google demand sources in the DA simulation to ensure that any increase in publisher revenues caused by DA is not driven by an increase in the total number of demand sources that the publisher accesses. The counterfactual in which the publisher does not displace an existing demand source is covered in Section VIII.C: with an appropriate floor, DA is a *risk-free revenue improvement* for the publisher.

289. I find that DA provides substantial benefits for publishers whether or not they would otherwise have included AdX in their waterfall. DA increases total publisher revenue from remnant inventory by at least [redacted]% compared to a counterfactual waterfall containing AdX. Gains in total publisher remnant revenue are even larger (at least [redacted]%) when compared to the baseline in which AdX does not otherwise participate in the waterfall. The match rate, which is the percentage of all impressions that are successfully sold by the publisher (to AdX or another demand source) also increased in both counterfactuals, from [redacted]% to [redacted]% in Counterfactual 1, and from [redacted]% to [redacted]% in Counterfactual 2. These results are summarized in Table 4 below.

Table 4: Change in Publisher Outcomes on Remnant Inventory After Moving from the Optimal Waterfall to DA[561]

[redacted]

290. Revenue effects naturally varied among publishers. For a complete picture of these effects, I plotted the revenues earned by each publisher with and without DA in Figure 9. In this figure, each point corresponds to a publisher. The horizontal axis is a publisher's simulated revenue under the Counterfactual 1 waterfall (in CPM) and the vertical axis is a

[561] The simulation results are output through code/parse_da_results.py in my supporting materials, with the logs saved in code/logs/parse_da_results.txt. The figures are saved in code/figures/, and the relevant files are prefixed by da_results.

publisher's revenue with DA enabled. The vast majority ([redacted]%) of the points in Figure 9 lie above the 45-degree line (plotted as a dashed line), indicating that most publishers in the simulations experienced an overall revenue increase from DA compared to the Counterfactual 1 baseline. For the majority of publishers, the gains are substantial, with the median publisher experiencing a [redacted]% increase in remnant revenue.

Figure 9: Publisher revenue increases from DA under Counterfactual 1

[redacted]

291. Different publishers also experienced differential effects of DA under the second counterfactual of my simulations. To describe the variations by publisher in Counterfactual 2, I plotted their revenues with and without DA in Figure 10. The interpretation of this figure is akin to that of Figure 9 above: the vast majority of the data

points ([redacted]%) lie above the 45-degree line (plotted). This means that, in my simulations, most publishers enjoyed overall revenue increases from DA compared to the Counterfactual 2 baseline, with the median publisher experiencing a [redacted]% increase in remnant revenue.

Figure 10: Publisher revenue increases from DA under Counterfactual 2

[redacted]

292. I provide more details on my data, modeling assumptions, simulation methodology, and results in the Technical Notes in Section XIII.B.

F. Responding to Plaintiffs' and Their Experts' Allegations

1. Plaintiffs' and Their Experts' Allegations Lack Historical Context and Ignore the Technical Challenge of Designing Unified Auctions

293. Plaintiffs compare DA to a counterfactual in which "exchanges compete at the same time for the impression by returning live, competitive bids,"[562] concluding that DA "ultimately reduced publishers' yield by shielding AdX from real-time competition."[563] Plaintiffs' experts also allege DA was "primarily motivated by Google's desire to maintain an information advantage over other exchanges," and that "Google's Dynamic Allocation [...] distorted the playing field in its favor because Dynamic Allocation was solely granted to AdX and not competing exchanges."[564]

294. These allegations fail in several ways to account for historical context.

295. *First,* DA was a significant improvement when evaluated in the context during which it was created. At the time of DA's launch, it was common for indirect sales channels to offer publishers *fixed* prices.[565] As noted by Professor Weinberg, "[i]nitially, all line items were static, so Dynamic Allocation addressed a natural shortcoming [...]."[566] When a publisher's indirect demand sources offer static prices, DA *must* increase publisher revenue and possibly increase the win-rate of remnant demand sources (for an example,

[562] Fourth Amended Complaint ¶ 270.

[563] Fourth Amended Complaint ¶ 274.

[564] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 549; Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 19.

[565] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -413 ("Ad networks provide revenue to publishers willing to commit, in advance, a large number of impressions at a fixed CPM value.").

[566] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 104.

see Paragraph 275).[567] Compared to the *status quo ante*, DA allowed real-time competition among advertisers, and according to Google's analysis, it *more than* [redacted] publishers' revenues on non-guaranteed impressions won by AdX.[568]

296. *Second,* in the period that followed the introduction of DA, incorporating real-time bids from additional indirect demand channels would still have required further technological progress, including agreements on applicable industry standards. When DA was introduced, ad exchanges were still nascent and the technology standard for real-time bidding had not yet been developed,[569] as Plaintiffs' experts have acknowledged.[570] At that time, industry participants viewed the main challenge not as exchange interoperability but "driving adoption" and having "sellers and buyers get comfortable with the exchange model."[571] DA eased this transition as it did not displace existing deals

[567] Suppose the average historical revenue from the best demand source is given by H and that the publisher sets a Value CPM, $V >= H$, for that demand. After DA, in the case where AdX clears V, the publisher must receive at least H. In the case where AdX does not clear V, the impression is allocated to the best alternative demand source, where the publisher receives at least H in expectation. In both cases, the publisher earns at least as much as they were earning before. Professor Weinberg acknowledges this outcome, noting that "[w]hen all other demand sources are static, Dynamic Allocation simply gives the publisher a shot at additional revenue." Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 112.

[568] White Paper, "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers" (2010), GOOG-DOJ-06818412, at -415 ("[A] publisher that generates an average CPM of $[redacted] for its non-guaranteed ad inventory might expect to see $[redacted] by taking advantage of the Ad Exchange.").

[569] Efforts to develop a standard began in November 2010 with the launch of the OpenRTB Consortium. *See* IAB Tech Lab, "OpenRTB" (Jul. 27, 2020), https://iabtechlab.com/standards/openrtb/ ("The Real-Time Bidding (RTB) Project, formerly known as the OpenRTB Consortium, assembled by technology leaders from both the Supply and Demand sides in November 2010 to develop a new API specification for companies interested in an open protocol for the automated trading of digital media across a broader range of platforms, devices, and advertising solutions.").

[570] *See* Expert Report of J. Gans (June 07, 2024), at ¶ 557 ("At the time that DoubleClick developed DA, the demand sources in the Waterfall were networks. Rather than running a real-time bidding auction and returning a live bid, networks simply purchased or did not purchase the impression when called in the Waterfall.").

[571] Email from [redacted] to [redacted], "FW: comments from industry players on AdX 2.0 on AdExchanger.com this evening" (Sep. 22, 2009), GOOG-AT-MDL-B-003180112, at -114 ("Regarding concerns, the biggest challenge for any exchange is driving adoption. It's not always as simple as 'if you build it, they will come'. The exchange represents a rather significant shift in how we typically transact, so adjusting to that for both buyer & seller takes some time."), -112 ("I think the launch of AdX 2.0 is an example of growing the pie as opposed to stealing share from a competitor because it's going to bring lots of new sellers (AdSense and DART for Publisher sites) and buyers (anyone who uses AdWords) into the market. As those sellers and buyers get comfortable with the

between publishers and remnant demand (including demand called by the waterfall), and this backward compatibility made it easy for publishers to adopt. Publishers could also disable DA if they decided it was not to their benefit.[572]

297. If it were even technically possible (which Plaintiffs' experts have not shown), implementing any unified auction in 2007 would also have been a significant organizational challenge, requiring demand sources to coordinate, agree on standards, and coordinate the times of implementation.[573] Such a process would have delayed the transition to a new ad allocation process. Redesigning Google's online display advertising platform to accept bids from other exchanges presented challenges, including the need to integrate bids in different auction formats, to avoid self-competition for bidders bidding on multiple exchanges, and to avoid price-fishing tactics by publishers. Responding to these challenges necessitated further innovations, as I discuss in Section XIII.

298. *Third,* Professor Gans claims that "DA allowed AdX, and only AdX, to compete in real-time against all non-guaranteed inventory, which was priced at a historical, average price, not a live auction price,"[574] but as Plaintiffs' experts and industry media

---

exchange model I think they'll begin trading on the other platforms as well, so I think this is a positive event for anyone in the exchange space.").

[572] *See, e.g.*, Summer Livestream Series (Sep. 2019), GOOG-AT-MDL-B-004582905, at -3162 ("You can still exclude specific ad units from dynamic allocation if desired."); "DFP and dynamic allocation" (Nov. 16, 2013), GOOG-DOJ-15416614, at -614 ("Publishers configure settings in DFP and Ad Exchange in order to control which inventory is eligible to compete, and how.").

[573] Efforts to develop a standard began in November 2010 with the launch of the OpenRTB Consortium, and support for header bidding was added to the OpenRTB standard in version 2.5, which was adopted in December 2016. *See* IAB Tech Lab, "OpenRTB" (Jul. 27, 2020), https://iabtechlab.com/standards/openrtb/ ("The Real-Time Bidding (RTB) Project, formerly known as the OpenRTB Consortium, assembled by technology leaders from both the Supply and Demand sides in November 2010 to develop a new API specification for companies interested in an open protocol for the automated trading of digital media across a broader range of platforms, devices, and advertising solutions. […] Open RTB 2.5 […] Release highlights include: […] Header Bidding Support - allowing for a signal when a bid request is originated from an upstream decisioning implementation like header bidding").

[574] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 568.

acknowledge,[575] the design of DFP did *not* prevent publishers or other exchanges from using a technology like header bidding to integrate real-time bids as early as the late 2000s.[576]

299. *Finally,* I note that other exchanges competed by offering a similar sales mechanism. For example, the OpenX ad server offered the OpenX exchange the opportunity to win an impression when it cleared the value CPM of the most competitive non-guaranteed line

[575] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 120 ("This is my opinion [...] accounting for periods both when other exchanges participated via the waterfall and when other exchanges participated via header bidding.").

[576] *See* Tim Cross, "Who Invented What in Ad Tech? – Part Two," VideoWeek (Jun. 26, 2018), https://videoweek.com/2018/06/26/who-invented-what-in-ad-tech-part-two/ ("Beeswax's Ari Paparo however backs up Brian O'Kelley's claim that he invented header bidding back in 2009, which AppNexus called 'pre-bid.' O'Kelley says the tech was a direct response to a DoubleClick for Publishers feature called 'Dynamic Allocation.' DFP, now owned by Google, allowed Google's own ad exchange AdX to compete against publishers' directly sold campaigns, while other exchanges were only able to compete for impressions not bought either through direct campaigns or AdX. Header bidding was developed as 'a hack around a feature gap in DFP,' says O'Kelley […] Header bidding's use spanned beyond simply countering DFP though, as it emerged as a popular alternative to the 'waterfall' system which offers impressions to demand sources one after the other […] Others however came upon the same idea through different route. Lee Cassingham said that his company Viewex came up with their own header bidding solution as part of a viewability product. Cassingham, having run various businesses within ad tech including an automated programmatic revenue/performance reporting tool, a viewability analytics tool, and an ad ops consultancy, looked for ways to create more value on a publishers' page, and created a 100 percent viewable ad format which appended itself to the bottom of a page. This then inspired a different solution to the viewability problem, where Viewex sought to build a capability for its own ad server whereby it could make viewability-based decisions on a web page before an ad loads. Whilst building the technology Cassingham realised he could capitalise on this pause by adding another step to call a third party to prospect for a higher bid and increase revenues […] Cassingham doesn't claim to have invented header bidding (Viewex's product came out in 2013), but his story provides an interesting example of how the same sort of technology can evolve independently while being built for completely different purposes.").

item,[577] [redacted][578]

2. Plaintiffs' and Their Experts' Allegations About the Purported "Right of First Refusal" Are Incorrect and Fail to Adequately Incorporate Incentives

300. Plaintiffs and their experts allege that DA afforded AdX a "right of first refusal" in which AdX could "peek at the average historical bids from rival exchanges and then transact the publisher's impression if AdX could return a live bid for just a penny more than the highest of these historical bids."[579]

301. This analysis is incorrect for several reasons. Whereas a typical "right of first refusal" grants a party the right to purchase an item under the same terms as a third party,[580] that is not the way that DA operates. Under DA, an AdX bidder needs to clear a value CPM chosen by the publisher in order to win an impression. That value CPM need not represent any other party's offer for that impression, so there is no "right of first refusal."

[577] *See* OpenX, "Selling rules" (Dec. 12, 2016), https://web.archive.org/web/20170708051049/https://docs.openx.com/Content/publishers/userguide_inventory_realtimeselling.html ("If enabled, you can designate inventory to sell through OpenX Ad Exchange using selling rules […] When OpenX receives an ad request for inventory defined by a selling rule, it proceeds with the selection process and selects an eligible line item for the ad space. If the selected line item is for non-guaranteed delivery, before serving the ad for the selected line item, OpenX will give buyers in the real-time bidding exchange the opportunity to bid on the ad space. If a bid is higher than the selected line item, and it matches or exceeds the floor price set for the selling rule, then OpenX serves the ad of the winning bidder, rather than the ad for the line item originally selected by the ad server.").

[578] *See* [redacted]

[579] Fourth Amended Complaint ¶ 271. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 548 ("AdX was afforded a right of first refusal on publisher's non-guaranteed inventory.").

[580] *See, e.g.*, Kahan, M., Leshem, S., & Sundaram, R. K. (2012). First-Purchase Rights: Rights of First Refusal and Rights of First Offer. American Law and Economics Review, 14(2), 331–371; Choi, A. H. (2009). A Rent Extraction Theory of Right of First Refusal. The Journal of Industrial Economics, 57(2), 252–264.

Furthermore, for inventory on which AdX did bid before other demand sources, the analyses of Plaintiffs and their experts fail to correctly incorporate incentives in two key ways.

302. *First,* in order to support their right-of-first-refusal interpretation, they presume incorrectly and repeatedly that real-time bids from AdX were compared to historical averages from other exchanges.[581] This claim is wrong. Publishers could enter *any* value—not just the exchange's historical average price—as a value CPM. For example, a publisher might configure DFP to assign a value CPM of $2 to a demand source, even though its average payment from the demand source was historically $1. If it did so, then it would not be the historical average price that AdX had to beat.

303. This possibility is not just theoretical: publishers had an *incentive* to set value CPMs *higher* than the historical average revenues of demand sources, and there is evidence that many publishers did just that.[582] Indeed, it is well known that it is *always* optimal for a seller to set an auction's floor price higher than the price it expects to outside the

[581] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 569 ("DA allowed AdX [...] to compete in real-time against all non-guaranteed inventory, which was priced at a historical, average price"). *See also* Fourth Amended Complaint ¶¶ 271 ("Google's Dynamic Allocation program instead had DFP permit AdX to peek at the average historical bids from rival exchanges and then transact the publisher's impression if AdX could return a live bid for just a penny more than the highest of these historical bids."), 274 ("Publishers ranked exchanges to reflect the historical average prices paid by each exchange."), 275 ("With Dynamic Allocation, Google used DFP to allow AdX to swoop in and buy inventory at just a penny more than the depressed average historical bids returned by non-Google exchanges to DFP.").

[582] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 11 ("Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons."). *See also* "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -506 ("We've anecdotally heard from some publishers that they inflate the value CPM of remnant line items to try and extract more value from AdX (since the remnant line item can set the reserve price for AdX 2P bids), to make it 'work harder.' [...] Comment [15]: This is based on anecdotal evidence (publisher conversations) - publishers used to do this even before HB was popular.").

auction.[583] Setting a floor price in that way increases the publisher's average revenues by forcing AdX bidders to pay more on average to win an impression.

304. Although Professor Weinberg confirms this incentive, he is dismissive about it, claiming that it would take a "sophisticated" publisher to recognize that it should set a higher floor price for the party getting a first-look at inventory.[584] That is wrong. For publishers, ad revenue is a significant income source, so they can be expected to pay attention to their revenue streams and employ sensible pricing strategies, guided either by analysis or by experimentation.

305. The basic pattern of optimal floor pricing in a sequential mechanism like Dynamic Allocation is not hard for a businessperson to understand. Just as the manager of a clothing store finds it best to set higher prices early in the season and lower prices later in the season to reduce the risk of unsold inventory, a publisher selling impressions finds it best to set a higher price for the first buyer in the waterfall and lower prices for later buyers for the same reason.

306. Even if a publisher sets suboptimal floor prices, the gains from Dynamic Allocation can still be large. In my simulations, when publishers set a floor price *equal* to the historical

---

[583] *See, e.g.*, Krishna, V. (2010). *Auction theory* (2nd ed.). Academic Press, at 23 ("[A] revenue maximizing seller should always set a reserve price that exceeds his or her value."). For the purpose of the AdX auction run under DA, the publisher's "value" for the impression is its expected revenue from sale of the impression outside of AdX.

[584] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 114 ("A sophisticated publisher could ignore Google's suggested formulas and set the Value CPMs however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default. Throughout the text, I use the term 'sophisticated publisher' with the following context in mind: Publishers can have varying levels of sophistication when attempting to optimize their revenue in the online display advertising ecosystem. On one end, a 'typical' publisher may set parameters according to their ad server's suggested text without developing a detailed understanding of how those parameters are used. At the other end, a 'sophisticated' publisher may fully digest all available documentation and aim to optimize parameters based on their use case, ignoring suggested text. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation.").

expected revenue of the waterfall instead of setting them at the optimal, higher level, revenues still increase: by [redacted]% in the first and [redacted]% in the second counterfactuals.[585]

307. Professor Gans justifies his conclusion that "[p]ublishers may incur losses due to AdX's prioritization in DA" with an example that ignores both the functioning of DA and publishers' incentives to set value CPMs higher than historical revenues.[586] In his example, the publisher sets a "CPM rate" of $5 for a non-Google demand source "based on historical performance," and the highest bid in AdX is $4.[587] He claims that "DA's preference for AdX will lead to the impression being served by AdX at $4 rather than passing through it to the alternative demand source."[588] But this is not how DA works: if the publisher sets a value CPM of $5 for the non-Google exchange, AdX would need to have a bidder willing to pay *more* than that amount in order for the impression to be allocated on AdX. Moreover, a revenue-maximizing publisher would set a value CPM *higher* than the historical average revenue of the non-Google demand source, so that, in practice, even a bid on AdX of $5 would likely be insufficient to win the impression.

308. *Second*, Plaintiffs' description incorrectly suggests that AdX adjusted its bids in response to what it learned from "peek[ing]" at the value CPMs setting the floor price.[589] At the

[585] The calculation of the revenue comparison can be found in code/parse_da_results.py of my supporting materials, with the logs saved in code/logs/parse_da_results.txt. In both counterfactuals, when floor prices were set in this manner, the total publisher revenue under DA reduced relative to the better-optimized floor prices described in Section XV.B.5.

[586] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 575.

[587] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 575 ("For instance, a publisher may set a CPM rate of $3 on AdX and $5 on other exchanges based on historical performance. The highest bid in AdX is $4, whereas the alternative exchange has a unique demand source, bidding $6 for the same impression.").

[588] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 575.

[589] Fourth Amended Complaint ¶ 271.

time that DA was adopted, AdX used a second-price auction, which is a bidder-truthful auction, as I have explained in Section III.C.3.a. That is, a bidder would be incentivized to bid its own value *regardless* of the floor price. In that case, the ability to observe the floor price does not benefit AdX bidders or affect publisher revenues.[590] Indeed, neither DV360 nor Google Ads ever adjusted their bids in response to the auction's floor price in the AdX second-price auction.[591]

309. Plaintiffs claim that AdX could "swoop in and buy inventory at just a penny more than the depressed average historical bids returned by non-Google exchanges to DFP,"[592] which is also misleading, as it misrepresents the operation of DA. Because AdX used a second-price auction, it only paid a "penny more" than the floor price in auctions when there was a single bidder above the floor. In many cases, when two or more AdX bidders cleared the floor, AdX's payment was set by the second-highest bid. In other words, competition in the internal AdX auction could lead to payments to publishers far exceeding the floor price, which is very different from a right of first refusal.

3. Professor Weinberg's Allegations About "Win Rate" and "Ad Quality" Do Not Hold Generally

310. Professor Weinberg acknowledges that DA has two opposing effects on the AdX win rate—namely, that "(a) AdX is always visited first in the waterfall, and (b) AdX's reserve

[590] The descriptions of Google's buy-side DRS, Bernanke and Poirot programs I have seen (and cited elsewhere in this document) suggest that Google's bidding tools did not adapt its bids to floor prices on AdX, as predicted by this theoretical result.

[591] *See* Declaration of N. Jayaram (May 1, 2024), GOOG-AT-MDL-C-000017969, at ¶ 3 ("Before Google Ad Manager transitioned to a Unified First-Price Auction in 2019, Google Ads and DV360 never used a floor price in an auction for an ad impression to adjust the bid values that they would submit for that same ad impression.").

[592] Fourth Amended Complaint ¶ 275

may be increased."[593] Despite this recognition, he reaches the unwarranted conclusion that DA "led to [a] higher win rate [...] for AdX,"[594] even though it is well understood that these opposing effects can result in a *lower* win rate for AdX. As a recent published economic analysis demonstrates, when an exchange such as AdX is called first, the publisher has an incentive to increase the exchange's floor price, sometimes by so much that its win rate *decreases*.[595]

311. In addition to his claims about win rates, Professor Weinberg reasons that, "***[i]f*** it is the case that AdX typically transacts ads that are of lower quality compared to non-Google Exchanges," then DA "would result in lower quality ads displayed on high-value impressions [...]."[596] Professor Weinberg does not provide evidence that ad quality differences between AdX and third-party demand sources were in fact common or significant, and, even if he did provide such evidence, his conclusion does not follow logically from his premise. For the sake of argument, suppose that "AdX typically transacts ads that are of lower quality."[597] In that case, the publisher using Dynamic Allocation could control the average ad quality in several ways, including by increasing the floor price applying to AdX (because doing so would increase the probability that the non-AdX bidder—assumed to be relatively higher quality—would win the impression). As a result, the publisher could determine for itself its optimal tradeoff between prices

[593] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 121.

[594] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 120.

[595] Despotakis, S., Ravi, R., & Sayedi, A. (2021). First-price auctions in online display advertising. *Journal of Marketing Research*, *58*(5), 888-907, at "Example 1" on p. 895.

[596] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 130 (emphasis added).

[597] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 130.

paid for impressions and ad quality, and depending on the publisher's preferences, the average ad quality could be higher or lower than in the absence of DA.

4. Plaintiffs' and Their Experts' Claims About "Cream-skimming" Are Overstated

312. Plaintiffs and their experts allege that DA permitted Google to "cream skim publishers' high-value impressions."[598] I interpret this claim to mean that DA allowed AdX to purchase impressions that would also have been valued highly by other demand sources, without those demand sources having a chance to bid.

313. Professor Weinberg describes this claim using an example, in which a well-matched advertiser, such as a running-shoe advertiser, bids in real-time through AdX and wins an impression for "a male over the age of 25 who likes running, and has visited ten different shoe company's [sic] websites in the last hour,"[599] over an advertiser bidding through a third-party demand source. There are several issues with his example and subsequent claims.

314. *First*, Professor Weinberg's example describes an atypical case in which all advertisers have high values for the same impression. In the more typical case, as acknowledged by Professor Weinberg,[600] bidders have heterogeneous values for an impression. For example, the end user in Professor Weinberg's example may have a variety of other

[598] Fourth Amended Complaint ¶ 281. *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 123 ("When other exchanges primarily participate via the waterfall, Dynamic Allocation, no matter how a publisher sets Value CPMs, would lead to AdX winning an even greater volume of high-value impressions, and increased revenues from these impressions under Dynamic Allocation compared to no Dynamic Allocation.").

[599] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 156.

[600] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...]. This is a simplifying assumption [...], and it is a sensible assumption to make since [...] bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's)").

interests that appeal to a broader pool of advertisers: he may have also recently browsed for flights, searched for a local gym membership, researched energy drinks, or searched for diapers for his baby daughter. In that case, the fact that the shoe advertiser bidding through a third-party demand source did not win the impression would more typically indicate that a bidder with a different interest in the end user won, in which case DA may help allocate the impressions to an advertiser with a higher value for that impression. By improving matching, DA increases revenues for publishers and may simultaneously increase the surplus of advertisers.

315. *Second*, Professor Weinberg suggests that publishers cannot use floor prices to effectively mitigate the issue in his example because "[AdX's] static reserve price is set based only on coarse targeting data [...]."[601] But Google Ad Manager does allow publishers to include rich targeting criteria in their line items, including "Custom targeting" that would permit the publisher to integrate live information about the end user's recent interactions with its website.[602] Even in cases where publishers may be unable to incorporate every bit of live information to update the static prices another bidder might offer, publishers could still raise the AdX floor price to offset any *expected* informational advantage that a bidder on AdX might have. Indeed, Google made this easier for publishers via its implementation of RPO, discussed in Section XI.

[601] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 123.

[602] Google, "Targeting types," Google Ad Manager Help (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/2884033?sjid=5613954590018599404-NC#custom-targeting ("Custom targeting allows you to include key-values, audience segments, or content metadata for video line items if Video Solutions is activated in your network. Key-values in particular can be used for purposes not captured by the built-in targeting in Ad Manager. For instance, you define key-values that identify specific ad inventory on web pages or apps, or they can be used to target ads based on that information you might gather from visitors to your website or app. Key-values, like other targeting, help your advertisers and buyers reach their intended audience or demographic.").

316. *Third,* Professor Weinberg's claims focus on the effects on competitors, rather than the effects on competition. After claiming that DA "would lead to AdX winning an even greater volume of high-value impressions,"[603] he does not elaborate as to how that would damage Google's own publisher-customers or advertisers or the functioning of the wider marketplace. DA has the pro-competitive effect of allowing Google to offer better services to its publisher-customers.

317. Even Professor Weinberg's conclusion that "other demand sources necessarily win a lower volume of these high-value impressions" focuses on the wrong measure.[604] He assesses the *average value* of the impressions won to measure impacts on non-Google demand sources, rather than the demand source's *total surplus*. Even if DA reduced the average value of impressions that a third-party demand source won, that demand source's total surplus might have increased if it transacted more impressions or faced lower floor prices due to its later position. As I demonstrated in an earlier example (see Paragraph 275), DA could make non-Google demand sources better off by increasing the volume and average value of impressions they would win.

[603] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 123.

[604] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 123.

IX. ENHANCED DYNAMIC ALLOCATION: INCREASING PUBLISHER REVENUES BY FULFILLING GUARANTEED CONTRACTS MORE EFFICIENTLY

A. Overview

318. Enhanced Dynamic Allocation (EDA) was a Google program to increase publishers' revenues and improve efficiency by allowing real-time bids to compete for impressions that, without EDA, would be allocated to guaranteed contracts. EDA expanded the pool of impressions for which remnant demand (including non-Google demand sources) could compete, dynamically allocating impressions between guaranteed campaigns and remnant demand to increase publisher revenues.[605]

319. At the time of launch, Google described EDA to publishers simply and accurately as a process "designed to increase your Ad Exchange revenue without compromising reservations" by assigning impressions to guaranteed line items "often enough to stay on pace to satisfy its goal."[606] EDA would assign an impression to remnant demand only if it would "pay more than the opportunity cost of not serving the guaranteed line item."[607]

---

[605] *See* Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -161 ("Enhanced Dynamic Allocation […] introduces competition between reservations and AdX by allowing AdX (or AdSense) to bid on high priority DFP impressions […] without compromising reservation goals. […] Enhanced Dynamic Allocation increases eCPMs and revenues without compromising reservations.").

[606] "EDA / AdX help articles for review: DFP and dynamic allocation" (Apr. 22, 2014), GOOG-DOJ-15417688, at -691 to -692.

[607] Google, "Delivery basics: Ad competition with dynamic allocation," Google Ad Manager Help (accessed Jan. 8, 2024), https://support.google.com/admanager/answer/3721872?hl=en.

320. Many other ad servers (including [redacted] Magnite, Comcast's Freewheel, [redacted], and OpenX) used technologies similar to EDA to allocate impressions between guaranteed contracts and remnant demand.[608]

321. Plaintiffs' and their experts' description of EDA contains various factual errors:

a. The complaint and Professor Gans assert that EDA offered AdX exclusive access to a new pool of publisher inventory.[609] However, EDA did not expand the pool for AdX bidders alone: it allowed *any* non-Google source of remnant demand

---

[608] *See* [redacted] Refinitiv Streetevents, "Q2 2021 Magnite Inc Earnings Call" (Aug. 5, 2021), https://investor.magnite.com/static-files/950ded3e-8953-4cf9-a043-d0e72b1fb856, at 4 ("Having a tightly integrated ad server allows for the dynamic allocation of programmatic and nonprogrammatic inventory to provide a holistic yield management solution for publishers."); OpenX, "OpenX transforms concept of SSP with industry-first demand fusion technology" (Jun. 9, 2014), https://www.openx.com/press-releases/openx-transforms-concept-of-ssp-with-industry-first-demand-fusion-technology/ ("Demand Fusion is a patent-pending technology that effectively fuses RTB and network demand, optimizing pricing through superior competition and maximizing revenue for publishers. This groundbreaking approach instantly evaluates all prices and dynamically selects the highest bid from both RTB and ad network buyers within the user's browser to optimize yield for each and every impression."); Ben Munson, "Comcast's FreeWheel launches unified ad decisioning," StreamTV Insider (Apr. 15, 2020), https://www.streamtvinsider.com/tech/comcast-s-freewheel-launches-unified-ad-decisioning ("Comcast's FreeWheel is launching a unified decisioning capability for buyers and sellers to transact across both direct sold and programmatic advertising.").

[609] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 630 ("As I explained above, EDA enables AdX (and only AdX) to transact impressions that would have been allocated to direct deals if it results in a higher clearing price. More specifically, AdX was given the ability to use the highest valued line item price as its reserve price, and transact the impression if it can beat this reserve price. No other exchange has this ability."). Fourth Amended Complaint ¶ 284 ("EDA had the purpose and effect of opening up a new additional pool of publishers' inventory to exactly one exchange: AdX.").

listed as a non-guaranteed line item to win impressions that would previously have been reserved for guaranteed deals.[610]

b. Plaintiffs suggest that EDA enabled Google to run auctions with a "[floor] price Google set for itself."[611] However, as Google told publishers, the EDA price was the outcome of a calculation designed to maximize publisher revenues while ensuring the delivery of publishers' guaranteed contracts. Moreover, EDA did not solely determine the floor price of the auction. EDA added an *additional* floor price, and publishers retained the ability to configure their own floor prices for the AdX auction.[612]

c. Plaintiffs' claim that "Google falsely told publishers that EDA 'maximizes yield,'"[613] is incorrect: EDA *did* maximize publishers' revenues.

322. Plaintiffs contend that EDA "cherry-picked" publishers' most valuable impressions away from direct sales channels.[614] Based on my analysis of Google data on impressions

[610] *See* Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -169 ("How does this affect other remnant line items? Remnant line items continue to function according to dynamic allocation rules - if a remnant [Line Item] (say, a network) has a higher eCPM than AdX, it will still get the impression with this optimization turned on."). *See also* Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 to -181.

[611] Fourth Amended Complaint ¶ 289 ("Google's exchange could transact the impression if an advertiser returned a net bid greater than both (a) the price Google set for itself and called the 'EDA reserve price' and (b) the average historical bids belonging to rival exchanges.")

[612] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 22 ("Between the launch of Enhanced Dynamic Allocation in 2014 and the launch of the Unified First Price Auction in 2019, the floor price in AdX was the highest of: (i) the publisher-configured floor price; (ii) the Enhanced Dynamic Allocation price set dynamically based on a temporary CPM (the 'EDA price'); (iii) the price of the remnant line item that was selected as a candidate for the impression; and (iv) the price determined by optimization.").

[613] Fourth Amended Complaint ¶ 291 ("Google automatically turned EDA on for publishers then coaxed publishers into leaving EDA turned on under a false pretense. Wearing its publisher ad server hat, Google falsely told publishers that EDA 'maximizes yield.' Publishers relied upon Google's misrepresentations to enable EDA, thinking it would maximize yield.").

[614] Fourth Amended Complaint ¶ 284 ("EDA had the purpose and effect of opening up a new additional pool of publishers' inventory to exactly one exchange: AdX. Moreover, this new pool contained publishers' most high-value

allocated to remnant demand and direct deals, I find no evidence of "cherry-picking" or "cream-skimming," as suggested by Plaintiffs. Although Professor Weinberg offers an example to support his opinion that "[t]he cream-skimming effect [...] potentially leads to a negative impact on the publisher in the long run,"[615] his example is unrepresentative, as it fails to incorporate the diversity of end user and advertiser interests. Instead, I expect that EDA's better matching would improve the experience for end users, increase revenues for publishers, and likely increase the surplus of advertisers.

B. How EDA Increased Efficiency and Publisher Revenues

1. EDA's Optimization Procedure

323. Prior to Google's introduction of Enhanced Dynamic Allocation, if an impression met the eligibility criteria for one or more guaranteed contracts, it would be immediately assigned to one of those. However, fulfilling guaranteed contracts in this way left money on the table for publishers. As a simple example, suppose that a user had visited Nike's online store, added a pair of shoes to their shopping cart and then left the site before purchasing the shoes. Despite Nike's high willingness-to-pay to show an ad to that user, DFP might assign the user's next impression to the publisher's guaranteed contract with 1-800-Flowers, even though that contract could be fulfilled later with a different impression. In general, a publisher has an incentive to sell an impression to remnant demand sources when they offer the highest price and to fulfill guaranteed contracts with

impressions (e.g., impressions displayed in the most prominent positions of a webpage, impressions targeted to users likely to make a purchase, etc.)."). *See also* Fourth Amended Complaint ¶ 292 ("Internally, Google understood that EDA was a scheme to let its own AdX exchange simply 'cherry-pick [publishers'] higher-revenue impressions,' earning Google's exchange an additional $150 million per year.").

[615] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 144 ("This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers. Hence, it potentially leads to a negative impact on the publisher revenue in the long run.").

the remaining impressions. The publisher does not, however, know in advance which impressions will fetch the highest offers from remnant demand sources.

324. In 2014, Google rolled out **Enhanced Dynamic Allocation (EDA)** to help publishers earn more revenue from remnant impressions by optimizing the selection of impressions used to satisfy guaranteed contracts.[616] At a high level, EDA worked as follows. After determining the best eligible guaranteed contract for an impression, the EDA algorithm calculated an **EDA price** to help allocate the impression.[617] Then, the EDA price served as an additional floor price for AdX and non-guaranteed line items.[618, 619] If no remnant buyer cleared the auction floor price, then the impression was assigned to the guaranteed contract. The EDA price was set to ensure that publisher revenues increased without jeopardizing fulfillment of the publisher's guaranteed contracts.[620] Indeed, by adding a floor price for non-guaranteed demand that was "informed by" the opportunity cost of not

[616] *See* Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -161 ("Generally [*sic*] availability roll-out began on 3/3/2014 […] Enhanced Dynamic Allocation […] introduces competition between reservations and AdX by allowing AdX (or AdSense) to bid on high priority DFP impressions […] without compromising reservation goals. […] Enhanced Dynamic Allocation increases eCPMs and revenues without compromising reservations.").

[617] Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 ("Enhanced Dynamic Allocation introduces competition between guaranteed reservations and other demand including AdX and DFP remnant reservations, by allowing AdX or DFP remnant to win over high priority DFP guaranteed reservations if it has a higher price than the *opportunity cost* (also called EDA price) set by us.").

[618] Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 ("Enhanced Dynamic Allocation introduces competition between guaranteed reservations and other demand including AdX and DFP remnant reservations, by allowing AdX or DFP remnant to win over high priority DFP guaranteed reservations if it has a higher price than the *opportunity cost* (also called EDA price) set by us. [...] DFP then ranks the remnant ads and picks one with highest eCPM. The selected remnant ad will be rejected if its eCPM is lower than the EDA price.").

[619] The EDA price was an additional floor price, and the effective floor price of the auction was set equal to the largest of the applicable floor prices (e.g. the EDA price, publisher-set floor price, etc.). Hence, in some instances, the effective floor could have exceeded the EDA price.

[620] Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 ("The EDA price is calculated in such a way that the DFP guaranteed reservation's delivery goal would not be compromised.").

serving a guaranteed line item, EDA worked similarly to the mechanism Professor Weinberg claims he would implement if he "were to design a waterfall-like format from scratch."[621]

325. To demonstrate how EDA's optimization procedure works, suppose that, over some period of time, a publisher expects to have 3,000 impressions to allocate, of which 2,000 must be assigned to guaranteed contracts, leaving 1,000 impressions for remnant demand. Suppose that, based on historical data, Google estimates that the publisher will receive bids of at least $1 from remnant demand on one-third of its impressions, and bids less than $1 on the remaining two-thirds. If these projections are accurate, the publisher would maximize its revenues on the 1,000 remnant impressions by only selling when the bid it receives is at least $1 and allocating the remainder of the impressions to the guaranteed contract. EDA achieves this by offering *each* impression to remnant demand with the EDA price of $1 as an additional floor price for the publisher. In that way, remnant demand purchases an impression only if it pays at least $1.[622] This results in the guaranteed contract being fulfilled and the set of impressions assigned to remnant demand being the ones with the highest bids from remnant bidders.

326. Implementation of this idea involved some subtle engineering details. One important detail arises because the bids that arrive over time can be forecast only imperfectly, so to

[621] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 76 ("Prioritization of direct deals over exchanges is a curious feature of the waterfall. If I were to design a waterfall-like format from scratch and I were unconstrained by technological challenges, I would (a) find the maximum payment v I could get from a direct deal for this impression (maybe v = 0, if it satisfies no direct deal targeting criteria), then (b) visit exchanges in the waterfall but setting reserves informed by v"). As Professor Weinberg notes, his suggestion is an idealized one which is "unconstrained by technological challenges." EDA works similarly, but was integrated as an iteration of Dynamic Allocation.

[622] *See* Section XV.C.1 for more detail on calculation of the EDA price, based on Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180-181.

ensure that all guaranteed contracts are fulfilled, AdX needed to track the rate at which guaranteed contracts were being served and update the EDA price adaptively.[623] Continuing the previous example, if the publisher's guaranteed contract was "behind schedule" at some point during the campaign (meaning that EDA had assigned to that contract fewer than two-thirds of the impressions received up until that point), then the EDA price would be increased above $1 and a larger fraction of the subsequent impressions would be assigned to the guaranteed contract. Conversely, if the guaranteed contract was "ahead of schedule" (meaning that EDA had assigned to that contract more than two-thirds of the impressions received up until that point), the EDA price would be reduced below $1. By carefully increasing and reducing the EDA price, Google could eliminate almost any chance that a guaranteed contract would not be fulfilled. As a second important detail, the proper handling of static bids sometimes required randomizing between serving the static bid or the guaranteed contract. I discuss this randomization in detail in the Technical Notes in Section XV.C.1.

2. EDA's Benefits for Publishers

327. Professor Weinberg emphasizes that EDA "led to an increase in win rate and increase in revenue for AdX"[624] and the complaint adds that "EDA hurt publishers' yield."[625]

[623] *See* Google, "DFP and dynamic allocation," DoubleClick for Publishers Help (captured on Sep. 22, 2015) at https://web.archive.org/web/20150922150140/https://support.google.com/dfp_premium/answer/3447903 ("The lower a line item's Satisfaction Index (SI) (that is, the more behind schedule it is), the higher the temporary CPM that's passed to Ad Exchange. Therefore, a standard line item that is behind schedule will win often enough to stay on pace to satisfy its goal and pacing settings.").

[624] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 101.

[625] Fourth Amended Complaint ¶ 289 ("Moreover, EDA hurt publishers' yield by permitting AdX to transact publishers' impressions for depressed prices. DFP permitted AdX to transact high-value impressions for one penny more than a price floor that Google set for itself—despite Google's obvious conflicts of interest. Google's exchange could transact the impression if an advertiser returned a net bid greater than both (a) the price Google set for itself and called the "EDA reserve price" and (b) the average historical bids belonging to rival exchanges.").

However, the introduction of EDA benefited publishers by increasing their revenues, so any change in AdX's win rate or revenue was competition on the merits: Google won more business by providing a better service to its publisher customers. This is formalized in the following theorem.

328. **Theorem 2:** Suppose that publishers' guaranteed contracts are unchanged after the introduction of EDA and that Google accurately forecasts the distribution of future bids from AdX. Then (1) EDA increases the publisher's expected revenue relative to the pre-EDA allocation procedure, and (2) if publishers set the optimal floor price for the AdX auction ignoring direct contracts, the floor set by EDA *maximizes* publisher revenue.

329. The proof of Theorem 2 is in Section XV.C.2. The intuition for why EDA increases publisher revenues is captured in the example in Paragraph 325 above. Under EDA, guaranteed contracts are still fulfilled, but remnant demand is no longer eligible to purchase an arbitrary selection of impressions—just those for which it is willing to pay most.

330. Moreover, if the EDA price is determined by an accurate forecast of the bids received by AdX and the publisher sets the optimal floor price ignoring direct contracts, the floor price chosen under EDA is optimal: if the EDA price was lower, there would not be enough impressions unsold on AdX to fulfill the guaranteed contract, while if the EDA price was higher (and higher than the publisher-set floor), then more impressions would be unsold on AdX than necessary to fulfill guaranteed contracts, costing publisher revenues. In practice, because Google cannot perfectly forecast the bids for all

impressions, EDA includes an adaptive procedure (discussed in Paragraph 112 above) to revise the floor price as bids are received during a campaign.

331. Google experiments confirm that EDA increased publisher revenues. Around the time of EDA's launch, a Google experiment found that "[t]ypically publishers see [redacted]% RPM [revenue-per-mille] increase, [with] some ad units seeing [a] [redacted]% increase."[626] In 2017, Google estimated that EDA delivered publishers "on average, a [redacted]%+ uplift in yield."[627] An experiment from 2017 found that EDA increased overall publisher revenue by nearly $[redacted] per day.[628]

3. How EDA Increased Efficiency

332. In addition to increasing publishers' revenue, EDA also improved efficiency by ensuring that fewer impressions went unsold and that the impressions allocated to remnant demand were the impressions for which remnant bidders had the highest values. Because EDA assigned to guaranteed contracts the impressions with the lowest bids from remnant demand (including the ones without bids that exceed the relevant floor price, which could not be sold to remnant demand at all), EDA could often reduce the number of unsold impressions, expanding output. Similarly, the total value of impressions allocated to remnant demand increased because the impressions allocated to remnant bidders were the ones for which they had the highest bids, which are also the impressions for which they had the highest values.

[626] Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -164 (under heading, "First Results").

[627] Presentation, "DoubleClick's Unified Stack - DFP + AdX" (Sep. 20, 2017), GOOG-DOJ-09163089, at -103.

[628] "Enhanced Dynamic Allocation" (Apr. 3, 2017), GOOG-DOJ-13208758, at -759 ("Publisher revenue increase is ~[redacted] / day or approximately [redacted] ARR.").

C. Responding to Plaintiffs' and Their Experts' Allegations

1. Plaintiffs and Their Experts Make Factual Errors in Their Description of EDA

333. The description of EDA by Plaintiffs and their experts contains several significant factual inaccuracies.

a. *First*, Plaintiffs state that EDA was "opening up a new additional pool of publishers' inventory to exactly one exchange: AdX,"[629] with Professor Gans asserting that "[n]o other exchange has this ability."[630] This is wrong. Whenever EDA allowed AdX to compete for an impression, it also allowed any non-guaranteed line items representing non-Google exchanges (and other demand sources) to compete.[631] This means that non-Google demand sources also benefited from EDA, allowing them to win additional impressions. An analysis of data from 2018-2019 found that around [redacted]% of the revenue increase caused by EDA came from remnant line items (rather than from buyers on AdX).[632]

b. *Second*, Plaintiffs assert that "Google's exchange could transact the impression if an advertiser returned a net bid greater than both (a) the price Google set for itself and called the 'EDA reserve price' and (b) the average historical bids belonging

629 Fourth Amended Complaint ¶ 284.

630 *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 630.

631 *See* Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -169 ("How does this affect other remnant line items? Remnant line items continue to function according to dynamic allocation rules - if a remnant [Line Item] (say, a network) has a higher eCPM than AdX, it will still get the impression with this optimization turned on.").

632 Emails from [redacted] to N. Korula, "Re: EDA Performance Remnant v.s. AdX" (Oct. 23, 2019), GOOG-AT-MDL-008935836, at -836 (Tables, column "extra remnant rev / extra total").

to rival exchanges."[633] Any suggestion that Google used EDA to manipulate the auction's reserve price in its own favor is incorrect. As Google told publishers, the EDA price was calculated to maximize publishers' revenues while ensuring the fulfillment of publishers' guaranteed contracts.

c. *Third*, Plaintiffs state that EDA "permitted AdX to transact high-value impressions for one penny more than a price floor that Google set for itself,"[634] but that description ignores the other floor prices set by publishers. EDA added an *additional* floor price, meaning that winning AdX bidders needed to exceed the floor price set by the publisher, in addition to the EDA price and other remnant bids.[635] Moreover, as I have explained in Section III.C.3.a, the fact that the winning bidder pays the larger of the floor price and the highest competing bid is the *defining* feature of the second-price auction format. In the second-price auction, the ability to observe the floor price (or a competitor's bid) offers no advantage to the bidder.

d. *Fourth*, Plaintiffs claim that "Google falsely told publishers that EDA 'maximizes yield.'"[636] However, as I demonstrate in Theorem 2, EDA *did* maximize the revenue from a publisher's demand sources. A lower EDA price would have resulted in a publishers' guaranteed contracts going unfulfilled, and a higher EDA

[633] Fourth Amended Complaint ¶ 289.

[634] Fourth Amended Complaint ¶ 289.

[635] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 22 ("Between the launch of Enhanced Dynamic Allocation in 2014 and the launch of the Unified First Price Auction in 2019, the floor price in AdX was the highest of: (i) the publisher-configured floor price; (ii) the Enhanced Dynamic Allocation price set dynamically based on a temporary CPM (the 'EDA price'); (iii) the price of the remnant line item that was selected as a candidate for the impression; and (iv) the price determined by optimization.").

[636] Fourth Amended Complaint ¶ 291.

price would have risked lower revenues from the sale of impression. Because the higher of the EDA price and the publisher-set floor applied to an impression, a publisher could set floor prices higher than the EDA price if they found it in their interest to do so.[637]

e. *Finally*, Plaintiffs argue that EDA allowed AdX to "transact publishers' impressions for depressed prices."[638] But by adding an additional floor price to a second-price auction, EDA only *increased* the prices at which publishers' impressions were sold. Moreover, Google's experiments found that EDA led to a significant increase in publisher's revenues from remnant demand, on the order of [redacted]%.[639]

2. EDA Does Not Leave Non-Google Exchanges with Lower Value Inventory

334. Plaintiffs and their experts allege that "EDA allowed AdX to transact a higher proportion of publishers' high-value inventory," resulting in "cherry-picking" or "cream-skimming" in which "more valuable transactions [are] transacted through AdX rather than direct deals."[640] Professor Weinberg adds that "[t]his cream-skimming effect reduces the value

---

[637] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 22 ("Between the launch of Enhanced Dynamic Allocation in 2014 and the launch of the Unified First Price Auction in 2019, the floor price in AdX was the highest of: (i) the publisher-configured floor price; (ii) the Enhanced Dynamic Allocation price set dynamically based on a temporary CPM (the 'EDA price'); (iii) the price of the remnant line item that was selected as a candidate for the impression; and (iv) the price determined by optimization.").

[638] Fourth Amended Complaint ¶ 289 ("Moreover, EDA hurt publishers' yield by permitting AdX to transact publishers' impressions for depressed prices. DFP permitted AdX to transact high-value impressions for one penny more than a price floor that Google set for itself—despite Google's obvious conflicts of interest. Google's exchange could transact the impression if an advertiser returned a net bid greater than both (a) the price Google set for itself and called the 'EDA reserve price' and (b) the average historical bids belonging to rival exchanges.").

[639] Comms Doc, "Enhanced Dynamic Allocation" (Nov. 6, 2017), GOOG-DOJ-06885161, at -164 (under heading, "First Results"); Presentation, "DoubleClick's Unified Stack - DFP + AdX" (Sep. 20, 2017), GOOG-DOJ-09163089, at -103.

[640] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 172 ("EDA allowed AdX to transact a higher proportion of publishers' high-value inventory by being able to bid for impressions previously reserved for direct deals."). *See*

of direct deals with publishers in the long run [and] potentially leads to a negative impact on the publisher revenue in the long run."[641]

335. To investigate the claim that EDA resulted in guaranteed contracts receiving "lower value" impressions, I performed an empirical analysis of Google data, using industry-standard quality measures to compare the quality of the impressions that EDA allocated to guaranteed contracts to the quality of all impressions that satisfied the contract terms.[642] The data I analyzed were from a time in which many non-Google exchanges participated in header bidding (see Section X) and/or Open Bidding (see Section XIII) and thus used real-time bids, allowing them to bid more for higher-quality impressions.

336. For the quality of each impression, the industry-standard measures that I use are **predicted engagement metrics**. When there are several guaranteed contracts eligible for an impression, Google uses these to decide which one to serve to maximize predicted engagement, that is, Google's estimate of the fraction of impressions for which a user

also Fourth Amended Complaint ¶ 292 ("Google understood that EDA was a scheme to let its own AdX exchange simply 'cherry-pick [publishers'] higher-revenue impressions.'"). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), from ¶ 142 to ¶ 144 ("If impressions that satisfy targeting criteria for direct deals are on average more valuable than impressions that do not, then Enhanced Dynamic Allocation results in more valuable transactions being transacted through AdX rather than direct deals. Therefore, Enhanced Dynamic Allocation would not only lead to increased volume and revenues for AdX, but also to a greater volume of valuable impressions being transacted through AdX[…]. This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers.").

[641] *See* Expert Report of M. Weinberg (Jun. 7, 2024), from ¶ 142 to ¶ 144 ("If impressions that satisfy targeting criteria for direct deals are on average more valuable than impressions that do not, then Enhanced Dynamic Allocation results in more valuable transactions being transacted through AdX rather than direct deals. Therefore, Enhanced Dynamic Allocation would not only lead to increased volume and revenues for AdX, but also to a greater volume of valuable impressions being transacted through AdX[…]. This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers.").

[642] The Google metrics I use below comply with industry standards for measuring clicks, views, and active views. *See* Media Rating Council, Minimum Standards for Media Rating Research, Media Rating Council (accessed Sep. 29, 2023), https://mediaratingcouncil.org/standards-and-guidelines.

engages with a particular ad. These metrics are the **predicted Click-Through Rate (pCTR)**,[643] **predicted View-Through Rate (pVTR)**,[644] and **predicted Active View Rate (pAVR)**.[645]

i) Data

337. I received a week of data, from October 30th to November 5th, 2022 corresponding to over [redacted] different guaranteed contracts.[646] The included contracts are all the ones that, on any day, (1) won and served at least 100 impressions under EDA and (2) for at least one of the three predicted-engagement statistics, lost at least 100 impressions under EDA.[647]

338. For each guaranteed contract, the dataset includes the numbers of impressions won and lost by the guaranteed contract and the predicted engagement metrics—pCTR, pVTR and pAVR—as described above. For each guaranteed contract and each engagement metric, I have access to:

[643] *See* Google, "Ad Manager report metrics," Google Ad Manager Help (accessed Sep. 29, 2023), https://support.google.com/admanager/table/7568664?hl=en ("Ad Exchange CTR[:] The percentage of impressions served by the Ad Exchange that resulted in users clicking on an ad.").

[644] *See* Google, "Understand Video Solutions report metrics," Google Ad Manager Help (accessed Jul. 24, 2024), https://support.google.com/admanager/answer/2759433#video-viewership ("View-through rate[:] Percentage of times the ad was viewed.").

[645] *See* Google, "How Active View metrics are calculated," Google Ad Manager Help (accessed Sep. 29, 2023), https://support.google.com/admanager/answer/6233478?hl=en&ref_topic=7506202 ("Active View metrics are calculated by finding the percentage of total impressions served that are actually viewable impressions").

[646] EDA pCTR Dataset, GOOG-AT-DOJ-DATA-000066771 to -772. The number of guaranteed contracts was calculated in the file code/pctr.py of my supporting materials. The output is logged to code/logs/pctr.txt.

[647] *See* Letter from D. Pearl to M. Freeman (Sep. 8, 2023), GOOG-AT-MDL-C-000012795, at -802 to -803 (describing the EDA pCTR dataset). For some contracts or impressions, Google does not compute any engagement metrics. Those impressions are not included in the dataset.

a. **Field #1:** The average value of the predicted-engagement metric over all impressions that would be eligible to fulfill the guaranteed contract.

b. **Field #2:** The average value of the predicted-engagement metric over the impressions allocated to the guaranteed contract.

ii) Methodology

339. I treat Field #1 as a measure of impression quality in a counterfactual "without EDA" scenario. I expect Field #1 (which contains the average value of the predicted-engagement metric over all eligible impressions) to be equal to the average value of that metric over impressions assigned to guaranteed contracts without EDA. The reason is that, before EDA, the impressions assigned to guaranteed contracts were the first eligible impressions to arrive in the relevant time period and, in terms of ad quality, I have no reason to believe that those impressions differ systematically from an average impression. Field #2 contains a measure of the quality of impressions that were allocated to guaranteed contracts in the actual world "with EDA." I compare values in the two fields to assess how EDA affected the quality delivered to guaranteed contracts. If adverse selection were an economically significant factor, the average measure of quality "without EDA" (Field #1) would be significantly larger than the average measure of quality "with EDA" (Field #2).

iii) Main Finding: Differences in Predicted-Engagement Metrics are Less than [redacted]%

340. I first computed the average predicted engagement metrics with and without EDA across all guaranteed contracts. I display my findings in Table 5.

Table 5: Change in Engagement Metric With EDA[648]

[redacted]

341. The main findings are:

a. *The* **predicted click through rates (pCTR)** *with and without EDA are nearly the same*. In terms of numbers of clicks, Google predicts [redacted] clicks per 100,000 impressions without EDA and [redacted] clicks with EDA. This means that the number of expected clicks for guaranteed contracts with EDA is about [redacted]% less than without EDA.

b. *The* **predicted view through rates (pVTR)** *with and without EDA are nearly the same.* Impressions allocated to guaranteed contracts via EDA have an approximately [redacted]% higher average pVTR, which means that the number of view-throughs for guaranteed contracts is about [redacted]% more than before EDA.

c. *The* **predicted active view rates (pAVR)** *with and without EDA are nearly the same.* The impressions won by guaranteed contracts under EDA have an approximately [redacted]% lower pAVR, which means that the number of active views for guaranteed contracts is about [redacted]% less than before EDA.

[648] These results were calculated in the file code/pctr.py of my supporting materials. The outputs are logged to code/logs/pctr.txt.

342. If the effects of "cream-skimming" were economically significant, one would expect the predicted engagement metrics on impressions allocated to guaranteed contracts to be substantially lower than the same metrics on the set of all eligible impressions. But they are not. Instead, the data shows decreases in just two of the three predicted engagement metrics, with all differences just [redacted].

iv) Publisher-Level Heterogeneity

343. The preceding results are average engagement rates over all publishers and guaranteed contracts. Averaging at this level still leaves open the possibility that some publishers could have suffered a significant loss. To evaluate that possibility, I conducted a similar analysis at the publisher level. I present my findings using a scatterplot in which each point represents a different publisher. On the two axes of the plot are the publisher's engagement metric in the EDA data vs. the same metric in the hypothetical pre-EDA scenario. Figure 12 shows the publisher-level scatterplot for pCTR.[649] If publishers experienced no selection effects, all points would lie exactly on the 45° line (the black dashed line in Figure 12), indicating that the publisher's pCTR is just the same with or without EDA. Points below the line correspond to publishers with worse pCTR impressions with EDA than without EDA ("adverse selection"), while the points above the line correspond to publishers with better impressions with EDA than without it ("positive selection"). The left panel shows the data for all publishers. The right panel is a zoomed-in version, focused at the range below 0.5% pCTR, where most observations lie. The colors of each point encode the logarithm of the number of impressions allocated to

[649] Figure 12 was generated by running the file code/pctr.py in my supporting materials. The figure is located at code/figures/pctr_regression.png.

guaranteed contracts in the data. I have also plotted (in red) the line of best fit to those points, using a weighted least-squares regression with zero intercept and with the number of winning impressions as weights.

Figure 12: Scatterplot of pCTR by publisher with/without EDA

344. Visually, the line of best fit (in red) and the 45° line (the dashed black line) are nearly indistinguishable, suggesting that EDA had a negligible effect on predicted engagement rates, uniformly across publishers. To express the same finding with a statistic, I calculate the slope coefficient of the best-fit line (the regression coefficient). If there was no difference in the predicted engagement rates with and without EDA, the regression coefficient would be 1.000. The computed regression coefficient is ____, suggesting that any difference in quality caused by EDA is very small indeed.

345. The computed regression coefficients for pAVR and pVTR are also suggestive of minor quality differences, with values of ____ and ____, respectively. The corresponding figures for those predicted engagement metrics can be found in Section XV.C.3 of the Technical Notes.

v) My Empirical Analysis Undermines Plaintiffs' Experts Conclusions About "High-Value" Ads

346. Together, these findings refute the idea that EDA systematically allocates "low-value" impressions to guaranteed contracts in the way that Plaintiffs and their experts suggest. On the contrary, there is barely any detectable quality difference between the impressions those contracts win or lose. The absence of a material, observed difference in quality in this data suggests that, on the dimensions of impression quality that Google uses for its own purposes, there is no economically important "cream-skimming." Moreover, since remnant bidders (including header bidders and exchanges in the waterfall) can condition their payments on more information than guaranteed contracts, it is reasonable to expect that, in all or most cases, non-Google bidders suffer even *less* harm than guaranteed contracts from any hypothetical "cream-skimming" by Google's ad-allocation processes, including DA and EDA.[650]

347. Instead of relying on data, Professor Weinberg forms his opinion on the basis of a hypothetical example, concluding that EDA reduces the quality of impressions served to guaranteed contracts and "*potentially* leads to a negative impact on the publisher revenue in the long run."[651] His example considers an "impression [] for a user over the age of 25 who likes running, [...] has recently visited several shoe-shopping websites, lives in a

[650] If "cream-skimming" was a significant factor in any part of the Google ad allocation process, then the residual impressions EDA allocates to guaranteed contracts should be of lower quality than the impressions allocated to either AdX or non-Google exchanges because these residual impressions attracted low bids—below the EDA price—not just by AdX bidders but also by competing exchanges.

[651] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 144 (emphasis added) ("This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers. Hence, it potentially leads to a negative impact on the publisher revenue in the long run.").

wealthy neighborhood, and competes in races."[652] The impression "satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of $10," but loses an auction to another advertiser (Reebok) bidding through AdX.[653] This is intended to show that EDA allows AdX bidders to "cream-skim" impressions that would be of equal or more value to guaranteed contracts.[654] However, Professor Weinberg's example is unrepresentative, as it oversimplifies the diversity of advertiser and end user interests.

348. End users are typically potential buyers for many different products on the web. For instance, the user in Professor Weinberg's example who likes running and races likely has other interests, too. Perhaps he may be planning a vacation or a move into a new home. Because users buy a variety of products, the same users are often targeted by advertisers from many industries, with advertisers differing in their values for each particular impression. In Professor Weinberg's words, "bidders are heterogeneous in their valuation for impressions."[655] By allowing a diverse pool of advertisers to compete for each user's attention, EDA leverages that bidder heterogeneity and facilitates matches that can be significantly better than when *only* the guaranteed advertiser can serve the impression. In

[652] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 135 ("Another impression arrives for a user over the age of 25 who likes running, and with further fine-grained cookies noting that the user has recently visited several shoe-shopping websites, lives in a wealthy neighborhood, and competes in races.").

[653] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 135 ("The highest bidder that exceeds their personalized reserve is Nike for $8 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of $10, which is the highest Value CPM among all line items. DFP then calls AdX with a reserve of $10, whose auction concludes with Reebok winning at a clearing price of $15. Reebok wins the impression through AdX, paying $15. This example is illustrated in Figure 27 below.").

[654] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 136 ("The last example also highlights one reason that AdX might possibly outbid high priority line items which typically have a much higher CPM compared to remnant line items, that the fine-grained targeting criteria may alert live demand sources of a particularly high value impression and bid higher than the direct deal price that is typically based on coarser targeting criteria.").

[655] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("[B]idders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's)[.]").

contrast to Professor Weinberg's example, when EDA allocates an impression to remnant demand instead of Altra's direct contract, that often happens because the user is targeted by an advertisement for a product *other than running shoes*. EDA can promote better matching by allocating the ad slot to Altra or Reebok when the user is researching running shoes, to American Airlines when he'd like to fly, or to a furniture seller during a move. The resulting matching tends to improve the experience for end users, increase revenues for publishers, and is also likely to increase the surplus of advertisers.

349. Furthermore, Professor Weinberg's example omits discussion of Altra's overall display advertising strategy, which might include a combination of guaranteed and programmatic advertising. In that case, EDA would make it *easier* for Altra to win valuable remnant inventory (including the remarketing impression described in Professor Weinberg's example) using real-time bidding. When Altra adopts this combination, it can effectively capture the attention of end users who have "visited several shoe-shopping websites" and also reach users who might not have been familiar with Altra. As a result, EDA can increase the value of impressions won by Altra on average across direct deals and remnant campaigns.

350. Professor Weinberg claims that, as a result of cream-skimming, EDA "reduces the value of direct deals with publishers in the long run[…] ***potentially*** lead[ing] to a negative impact on[…] publisher revenue in the long run."[656] However, Professor Weinberg's argument provides no actual evidence that EDA reduced the value of direct deal advertising or that publisher revenues were reduced. My analysis of the

[656] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 144 (emphasis added) ("This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers. Hence, it potentially leads to a negative impact on the publisher revenue in the long run.").

predicted-engagement data shows that, to the contrary, EDA did not reduce the value of direct deal advertising. But even if it were true that EDA "reduce[d] the value of direct deals,"[657] Professor Weinberg provides no justification for why any decrease in guaranteed contract revenue would outweigh the potentially significant increases in revenue from remnant demand sources. Indeed, if EDA increased the average revenue per impression allocated to remnant advertising, that would increase the competition for all impressions and contribute to an increase in the price of direct deals as well, which is the opposite of what Professor Weinberg's analysis implies.

351. Professor Pathak claims that "Google reduced publisher choice when it restricted publishers' ability to opt out of Enhanced Dynamic Allocation" and that "EDA diminished publishers' control […] for their high-value inventory."[658] But, as I have shown, EDA *maximizes* publisher revenues, while having no economically significant effect on the quality of impressions allocated to guaranteed contracts. Because EDA *increased* ad server functionality by unifying the competition between guaranteed and remnant demand, publishers had no incentive to "opt out."

---

[657] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 144.

[658] Expert Report of P. Pathak (Jun. 7, 2024), at ¶¶ 172 ("Google reduced publisher choice when it restricted publishers' ability to opt out of Enhanced Dynamic Allocation (or 'EDA') while using Google's advertising exchange AdX."), 173 ("EDA diminished publishers' control over demand sources for their high-value inventory. EDA allowed AdX to transact a higher proportion of publishers' high-value inventory by being able to bid for impressions previously reserved for direct deals. Publishers did not have the option to selectively turn off EDA for a selection of their premium inventory available through direct deals, and instead had to allow AdX the ability to bid for all of their premium inventory. Google therefore diminished publishers' ability to be selective with demand sources and publishers had to allow for the possibility of AdX serving lower-quality ad impressions for their most premium inventory.").

X. HEADER BIDDING AND THE DUBIOUS "LAST LOOK ADVANTAGE"

A. Overview

352. Header bidding is a technology that allows publishers to solicit and compare real-time bids from their selected sets of exchanges and other demand sources.[659] Typically, header bidding used a first-price auction to compare bids from the publisher's chosen demand partners.[660] If a publisher configured its website to run header bidding before offering an impression to AdX, AdX would obtain a so-called "**last look**" at each impression. The so-called "last look" was not a Google program, but arose as a consequence of the way that some publishers configured header bidding.[661] Those publishers benefited from offering AdX bidders the chance to bid on inventory (rather than using header bidding alone) because the additional competition for each impression allowed them to earn higher revenues and because it allowed them to enjoy the benefits of the services that GAM provided on each impression.

353. AdX was not the only exchange that collected bids after header bidding auctions had been resolved: other exchanges with ad servers, such as OpenX, treated header bids

[659] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("When a user visits the publisher's site, the browser calls participating ad exchanges or other demand partners (either directly or via a header bidding server) to submit bids, and runs an auction between those bids before Google's ad server is called.").

[660] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("These Header Bidding auctions are typically first-price auctions.").

[661] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 17 ("'[L]ast look' was not designed to give AdX an advantage when competing against Header Bidding. It was simply the result of the Header Bidding auction taking place before the AdX auction ran and the way that publishers configured Header Bidding line items to work with Dynamic Allocation.").

similarly.[662] As I discuss in Section XIV, when AdX transitioned to the Unified First Price Auction in 2019, AdX bidders and remnant line items (including those representing header bidding) competed simultaneously, eliminating any claimed "last look."

354. Plaintiffs allege that the so-called "last look" gave AdX an advantage over third-party exchanges because it "let AdX peek at the winning net bid from an exchange using header bidding, then displace the trade by paying one penny more."[663] However, these allegations are based on flawed analyses that systematically ignore how incentives affect bidder and publisher behavior.

a. *Bidder incentives*. While Plaintiffs analogize "last look" to "insider trading" and "front running"[664] and their experts claim that it "creates information asymmetries,"[665] the ability to see competing bids in a second-price auction (such

[662] OpenX, "Selling Rules" (Dec. 12, 2016), https://web.archive.org/web/20170708051049/https://docs.openx.com/Content/publishers/userguide_inventory_realtimeselling.html ("If the selected line item is for non-guaranteed delivery, before serving the ad for the selected line item, OpenX will give buyers in the real-time bidding exchange the opportunity to bid on the ad space. If a bid is higher than the selected line item, and it matches or exceeds the floor price set for the selling rule, then OpenX serves the ad of the winning bidder, rather than the ad for the line item originally selected by the ad server.").

[663] *See* Fourth Amended Complaint ¶ 377 ("From the earliest days of header bidding, DFP let AdX peek at the winning net bid from an exchange using header bidding, then displace the trade by paying one penny more. Industry participants called this practice, along with Dynamic Allocation, Google's 'Last Look.' Other industries call analogous conduct by intermediaries 'insider trading' and 'front running.' According to a confidential Google study evaluating the effects on competition, Last Look significantly re-routed trading from non-Google exchanges to AdX and Google's ad buying tools, protecting Google's market power in both. Google itself admitted: 'Last Look is inherently unfair.'").

[664] *See* Fourth Amended Complaint ¶ 377 ("Other industries call analogous conduct by intermediaries 'insider trading' and 'front running.'").

[665] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 594 ("Notably, after the introduction of Header Bidding, DFP could pass the winning Header Bidding bid to AdX as the price floor in its real-time auction. This advantage became known as 'Last Look', where AdX 'gets to bid with knowledge of the clearing price.'"). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 119 ("Hence, Dynamic Allocation allows AdX (and only AdX) to learn others' bids in a first-price auction format, and as a result, Dynamic Allocation creates information asymmetries that favor Google's AdX. This advantage is often referred to as AdX's Last Look advantage."). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 126 ("[U]nder Dynamic Allocation, AdX is the only exchange that learns information about others' bids and passes it on to its bidders[...] I previously noted that this is referred to as AdX's Last Look advantage and constitutes a significant advantage in an auction[...].").

as the one AdX used before 2019) does not confer an advantage to the bidder.[666] Indeed, bidders in a second-price auction are incentivized to bid their values regardless of competing bids or the floor price (as described in Section III.C.3.a). Sometimes a winning bid results in trade at the floor price, but—in those cases—*no* alternative bid would lead to a higher price for the publisher. Moreover, Google's own demand sources, Google Ads and DV360, never used bids from header bidding (or the value CPMs of header bidding line items) to adjust the amount they bid on the same impression.[667]

b. *Publisher incentives*. Publishers could maximize their revenues by inflating the header bids that they submit to AdX (that is, triggering a line item with a *higher* value CPM than the winning header bid), leading to higher floor prices on AdX.[668] Configuring header bidding this way would require a bidder on AdX to bid *more than* a penny above the highest bid made in header bidding in order to win. These higher floor prices reduce or eliminate any advantage that might arise as a result of header bidding being conducted in first-price format (in which bidders optimally bid less than their values, as described in Section III.C.3.b) whereas the AdX auction was conducted in second-price format. In practice, internal Google documents suggest that many publishers, in line with these economic incentives,

[666] Note that bidders on AdX (and AdX itself) did not observe header bids directly (see Paragraph 357).

[667] Declaration of N. Jayaram (May 1, 2024), GOOG-AT-MDL-C-000017969, at ¶ 3 ("Before Google Ad Manager transitioned to a Unified First-Price Auction in 2019, Google Ads and DV360 never used a floor price in an auction for an ad impression to adjust the bid values that they would submit for that same ad impression.").

[668] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 11 ("Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons.").

configured header bidding to trigger line items that set the AdX second-price auction floor *greater* than the highest header bid.[669]

355. Since 2014, Google's ad server has limited the maximum number of active line items to 61,000[670] in order "to protect the health of the product [and] the performance of [Google's] system."[671] Because some publishers implemented header bidding in GAM using many thousands of line items, Plaintiffs assert that this "throttles publishers' use of header bidding by artificially capping publishers [*sic*] use of 'line items' [...]."[672] However, my analysis of GAM data in Section X.D.4 below finds that publishers rarely approach the 61,000 threshold in practice. Although Professor Gans opines that "Google's motive in imposing limits to the number of line items… was to restrict

[669] *See* Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -506 ("[Some publishers] inflate the value CPM of remnant line items [...] publishers used to do this even before HB was popular."); Presentation, "First-price bidding" (Aug. 12, 2019), GOOG-DOJ-11406673, at -677 ("How boost works[:] The publisher inflates the HB bid before sending it as a floor to AdX[.] This is done to increase Adwords cost and to provide a better comparison between Adwords and header bidder bids[.]"). *See also* Email from [redacted] to [redacted] et al., "Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019" (Aug. 13, 2019), GOOG-DOJ-09713317, at -319 ("Today, these [publisher-]inflated CPMs are used to provide price pressure for AdX [...] In practice, [...] many publishers [...] apply a boost to Header Bidding bids[.]"); Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -777 ("Last look [...] incentivizes pubs to inflate ('boost') the floor sent to AdX"); [redacted]

[670] Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -198 ("Number of active line item limits (please note: Each creative-level targeting criterion counts also toward the active line item limit[:] 61,000").

[671] Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -195 ("In 2014, we have started enforcing limits on the creation of certain entities in the DFP Ul, e.g. Lls per order, Creatives per LI, Ad units for placements and Targeting attributes in line items. [...] Limits are necessary in order to protect the health of the product, the performance of our system, and are ultimately for the benefit of all publishers and the performance of their Ul (unlimited number of entities can affect the Ul in terms of load time and serving in some cases). NOTE: Our limits are much higher than those that existed previously in DART.").

[672] Fourth Amended Complaint ¶ 389.

competition in the ad exchange market" and dismisses Google's reasons for the limits as "pretextual,"[673] the documents he relies upon undermine his opinion and suggest that Google had legitimate technical reasons to protect its systems from the strain caused by publisher configurations with unusually many line items.

B. How the So-Called "Last Look" Arose As a Consequence of Some Publishers' Configuration of Header Bidding in GAM

356. **Header bidding** is a technology that allows publishers to request and compare real-time bids from multiple demand sources simultaneously.[674] It began to gain popularity around 2014,[675] years after Dynamic Allocation launched, as a way for publishers to incorporate real-time bids from multiple exchanges.[676] Publishers typically implemented header bidding through snippets of code (in the header of the web page) that sent bid requests from the end user's browser to ad exchanges.[677] Header bidding can be described as an

[673] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 646 ("Google's motive in imposing limits to the number of line items available to publishers, and denying requests for those limits to be raised, was to restrict competition in the ad exchange market by making Header Bidding more difficult and costly to the largest and most important publishers. While Google offered various technical explanations for the caps, these were pretextual.").

[674] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("When a user visits the publisher's site, the browser calls participating ad exchanges or other demand partners (either directly or via a header bidding server) to submit bids, and runs an auction between those bids before Google's ad server is called.").

[675] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("Around 2014, web publishers began to adopt Header Bidding."). *See also* AdPushup, "Header Bidding" (2023), https://www.adpushup.com/header-bidding-guide/ ("Header bidding made it to ad tech somewhere around 2014. And only after one year, in 2015, the technique went viral.").

[676] An internal document describes header bidding as a process motivated by publishers "clamouring" for "per-query bid" functionality from other exchanges, see Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337 ("Turns out that getting per-query bids from exchanges dramatically increases yield (auctions more efficient), so pubs are clamouring for this functionality") (emphasis omitted).

[677] My understanding of the configuration of header bidding in DFP is based on Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -338 to -339.

auction of auctions because demand sources would typically determine their bids using an internal auction before participating in the header bidding first-price auction.[678]

357. As described in Section III.D.3.f, header bidding could directly allocate the impression to the highest bidding exchange,[679] but, for the reasons I discuss in Section X.C below, many publishers sought to integrate header bidding with Google's ad server.[680] Because line items in DFP were not designed to accept real-time bids from non-Google exchanges,[681] these publishers used non-guaranteed line items—originally intended to represent static bids from ad networks—to instead represent real-time bids from non-Google exchanges.[682] To accomplish that, a publisher would configure its header bidding code to select a non-guaranteed line item with an accompanying value CPM in Google's ad server. These value CPMs could be chosen freely by the publisher and, importantly, could differ from the winning header bid. Prebid.js, one of the leading header bidding tools,

[678] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 13 ("These Header Bidding auctions are typically first-price auctions."). *See also* Prebid.org, "Prebid.js FAQ" (accessed Dec. 1, 2023), https://docs.prebid.org/dev-docs/faq.html ("Header Bidding is a first-price auction").

[679] *See* Prebid.org, "Running Prebid.js without an ad server" (accessed Sep. 7, 2023), https://docs.prebid.org/dev-docs/examples/no-adserver.html ("This example demonstrates running [a header bidding] auction and rendering without an ad server.").

[680] *See* Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -939 ("43% of LPS publishers are using header tags").

[681] Deposition of [redacted] at 71:18-21 (Aug. 12, 2021), GOOG-AT-MDL-007178292, at -363 ("So in the original design of the system, it was not designed to put exchanges in as line items. Line items are designed to represent direct deals or network deals."); Deposition of [redacted] at 50:8-20 (Nov. 6, 2020), GOOG-AT-MDL-007172126, at -176 ("The way the system was built is that line items were always intended to be reservations. There wasn't a concept of using them for realtime pricing. And so we had in mind that publishers would have, you know, possibly thousands of line items and the system was built to scale to that, but with using line items for realtime pricing, which is not what they were designed to be used for, there were ten, sometimes ten times, sometimes 100 times, sometimes 1,000 times more line items than the system was designed to support.").

[682] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 14 ("Up until at least December 2021, the winning bid from the Header Bidding auction was typically used to trigger a specific line item that the publisher had booked within Google's ad server (most commonly a remnant line item), and […] the Value CPM of that line item could represent the winning Header Bidding bid as a floor in the AdX auction (prior to September 2019) or as a competing bid in the Unified First Price Auction (from September 2019 onwards).").

contains an inbuilt feature to adjust header bids before they are sent to the ad server, with Pubstack, a supply-side tool, detailing how the feature could be used to "giv[e] Prebid an edge in the competition against GAM" by inflating header bids.[683] The actual header bids are never observed by GAM or bidders on AdX. GAM would use the value CPM of the header bidding line item in its ad selection process (including DA and EDA). As a result, the highest bid from an AdX bidder would win the impression only if it exceeded the highest value CPM associated with header bidding line items and any other floor prices that might apply to the impression. If there was no such higher AdX bid and the header bidding line item had the highest value CPM among all the other line items, the impression would be allocated to the winning header bidder.

358. In each auction, AdX bidders are informed of the highest floor price applying to that impression.[684] Because a header bidding line item sometimes determines the auction's floor price, critics have sometimes described AdX as having had a "last look" over header bids.[685] However, as just described and in further detail in Section X.D.1 below, the resulting floor price need not be *equal to* the winning header bid. In fact, publishers had an incentive to, and often did, set the floor price *higher than* the winning header bid,

[683] Asmaâ Bentahar, "Bid Adjustments Simplified: Run Fair Auctions with no Hassle," Pubstack (May 2, 2021), https://www.pubstack.io/topics/bid-adjustments-simplified ("Alternatively, this next one will slightly increase all returned CPMs, giving Prebid an edge in the competition against GAM").

[684] Before 2016, Authorized Buyers (then "AdX buyers") could only see the highest floor price among the "rule" floor prices set by the publisher and any floor price determined by RPO, but since 2016, all bidders see the highest among all the floor prices, including those determined by the value CPMs of remnant line items. *See* Launch Doc, "Including Third-Party Threshold in the Revealed Reserve Prices to AdX Buyers" (Aug. 9, 2016), GOOG-DOJ-13208800, at -800 ("We propose to include third-party threshold in the revealed reserve prices sent to AdX buyers, in support to AdX dynamic revshare v2 launches. Currently we reveal the maximum of rule prices and RPO prices in the RTB as well as JEDI callouts and in the internal RPC to GDN / DBM. With this change we will reveal the maximum of rule prices, RPO prices, as well as third-party threshold (e.g., EDA prices, competing DFP line item prices).").

[685] Presentation, "Changes to Ad Manager auction" (Jan. 10, 2019), GOOG-DOJ-10924270, at -273 ("AdX and EB have visibility into remnant price before they submit bids (commonly referred to in market as 'last look')").

which might allow the header bid to win even when the AdX bid was higher. By setting floor prices in this manner, a bidder on AdX would need to bid *more than* a penny above the highest bid from a header bidding exchange in order to win the impression.

359. The so-called "last look" was not a Google program: it arose as a consequence of the way that publishers configured header bidding using the line item capabilities that publisher ad servers (including DFP) supported at the time header bidding was introduced.[686] A system that incorporated non-Google exchanges reporting real-time bids directly to GAM required a major redesign, which Google later initiated with the implementation of Open Bidding and completed with the Unified First Price Auction, which I discuss further in Section XIII. Publishers using ad servers besides DFP employed similar techniques to integrate header bidding. Documentation from OpenX—another company that operated both an ad server[687] and an ad exchange—indicates that publishers using their ad server and exchange integrated header bidding in a similar way to publishers using GAM: OpenX bidders competed *after* third-party remnant demand, with header bids

[686] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 17 ("'[L]ast look' was not designed to give AdX an advantage when competing against Header Bidding. It was simply the result of the Header Bidding auction taking place before the AdX auction ran and the way that publishers configured Header Bidding line items to work with Dynamic Allocation.").

[687] OpenX shut down its ad server after 2018. *See* Sarah Sluis, "OpenX Lays Off 100 Employees And Pivots To Video," AdExchanger (Dec. 18, 2018), https://www.adexchanger.com/platforms/openx-lays-off-100-employees-and-pivots-to-video/.

determining floor prices.[688] [redacted].[689]

C. Publishers Benefited from the So-Called "Last Look"

360. Publishers could use header bidding to sell online display advertising impressions without ever calling GAM or another publisher ad server.[690] But offering AdX bidders a chance to bid on inventory provided two important benefits for publishers.

361. *First*, by allowing competition from an additional source of demand, publishers could increase their revenues on each impression. As long as publishers triggered a header bidding line item with a value CPM at least as large as the best header bid (which they had a clear incentive to do), GAM's ad selection process could increase revenue from the impression, but could never reduce it.

362. *Second*, by offering AdX bidders the chance to bid on an impression, publishers could take advantage of the other services provided by GAM. For example, a publisher that used GAM to manage guaranteed contracts could combine the benefits of header bidding and EDA: a publisher could run header bidding on an impression before calling GAM,

[688] OpenX treated header bids as non-guaranteed line items, *see* OpenX, "Selling Rules" (Dec. 12, 2016), https://web.archive.org/web/20170708051049/https://docs.openx.com/Content/publishers/userguide_inventory_realtimeselling.html ("If enabled, you can designate inventory to sell through OpenX Ad Exchange using selling rules […] When OpenX receives an ad request for inventory defined by a selling rule, it proceeds with the selection process and selects an eligible line item for the ad space. If the selected line item is for non-guaranteed delivery, before serving the ad for the selected line item, OpenX will give buyers in the real-time bidding exchange the opportunity to bid on the ad space. If a bid is higher than the selected line item, and it matches or exceeds the floor price set for the selling rule, then OpenX serves the ad of the winning bidder, rather than the ad for the line item originally selected by the ad server.").

[689] [redacted]

[690] Prebid.org, "Running Prebid.js without an ad server" (accessed Sep. 7, 2023), https://docs.prebid.org/dev-docs/examples/no-adserver.html ("This example demonstrates running [a header bidding] auction and rendering without an ad server.").

and EDA could allocate the impression to the guaranteed contract or the header bid (or another line item), depending on which was expected to maximize publisher revenues.

D. Responding to Plaintiffs' and Their Experts' Allegations

1. Plaintiffs and Their Experts Mistakenly Claim that AdX Buyers Benefit from Additional Information Provided by the So-Called Last Look

363. Plaintiffs allege that the so-called "last look" gave AdX's bidders the ability to "peek" at header bids, analogizing Google's alleged conduct to "insider trading" or "front running."[691] Plaintiffs' experts seemingly agree, with Professor Gans characterizing the so-called "last look" as an "advantage" because AdX "gets to bid with knowledge of the [Header Bidding] clearing price,"[692] and Professor Weinberg arguing that "Dynamic Allocation allows AdX (and only AdX) to learn others' bids in a first-price auction format, and as a result, Dynamic Allocation creates information asymmetries that favor Google's AdX. This advantage is often referred to as AdX's *Last Look advantage.*"[693] These allegations are completely wrong.

364. AdX ran a *second-price* auction at the time that DA was launched. As I have explained in Section III.C.3.a, observing others' bids does not advantage bidders in a second-price

[691] *See* Fourth Amended Complaint ¶ 377 ("From the earliest days of header bidding, DFP let AdX peek at the winning net bid from an exchange using header bidding, then displace the trade by paying one penny more. Industry participants called this practice, along with Dynamic Allocation, Google's 'Last Look.' Other industries call analogous conduct by intermediaries 'insider trading' and 'front running.' According to a confidential Google study evaluating the effects on competition, Last Look significantly re-routed trading from non-Google exchanges to AdX and Google's ad buying tools, protecting Google's market power in both. Google itself admitted: 'Last Look is inherently unfair.'").

[692] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 594.

[693] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 119. *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 126 ("[U]nder Dynamic Allocation, AdX is the only exchange that learns information about others' bids and passes it on to its bidders[...] I previously noted that this is referred to as AdX's Last Look advantage and constitutes a significant advantage in an auction.").

auction. The reason is that a surplus-maximizing bidder in a second-price auction bids *without regard to the floor price and the bids of other bidders*. The bidding programs used by Google's own demand sources in the AdX second-price auction conform to this theoretical prediction: none adjusted bids depending on the floor price of the auction.[694] If bids do not change in response to floor prices, the ability to see floor prices cannot harm publishers' revenues or benefit AdX bidders.

2. Plaintiffs and their Experts Neglect Publishers' Ability and Incentive to Inflate the Winning Header Bid

365. Plaintiffs allege that "last look" unfairly "re-routed trading from non-Google exchanges to AdX" by letting AdX "displace their trades by a penny."[695] Plaintiffs' experts echo this argument, with Professor Weinberg stating that "Dynamic Allocation enabled AdX to win impressions by bidding one cent above the header bidding clearing price"[696] and Professor Gans asserting that "last look increased AdX's probability of winning the impression."[697]

366. I interpret these allegations to mean that "last look" made an AdX bidder more likely to win (or pay less on average for) an impression than an otherwise similar bidder on a

[694] *See* Declaration of N. Jayaram (May 1, 2024), GOOG-AT-MDL-C-000017969, at ¶ 3 ("Before Google Ad Manager transitioned to a Unified First-Price Auction in 2019, Google Ads and DV360 never used a floor price in an auction for an ad impression to adjust the bid values that they would submit for that same ad impression.").

[695] *See* Fourth Amended Complaint ¶¶ 377-78 ("From the earliest days of header bidding, DFP let AdX peek at the winning net bid from an exchange using header bidding, then displace the trade by paying one penny more. Industry participants called this practice, along with Dynamic Allocation, Google's 'Last Look.' Other industries call analogous conduct by intermediaries 'insider trading' and 'front running.' According to a confidential Google study evaluating the effects on competition, Last Look significantly re-routed trading from non-Google exchanges to AdX and Google's ad buying tools, protecting Google's market power in both. Google itself admitted: 'Last Look is inherently unfair.' [...] These exchanges could now also peek at header bidding net bids and displace their trades by a penny.").

[696] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at Section IV.B.2.b.

[697] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 548.

non-Google exchange, even when those bidders have the same value for the impression.[698] Plaintiffs' and their experts' conclusions are based on an incorrect qualitative analysis: after properly accounting for bidder and publisher incentives, my economic analysis finds that the so-called "last look" did not provide AdX with an inherent advantage.

367. Plaintiffs' arguments ignore how publisher incentives affect their choices of floor prices. The claim that the so-called "last look" allowed AdX bidders to pay "a penny more" than the header bid tacitly assumes that AdX bidders could see the winning header bid as a floor price in AdX, but this is incorrect because any floor price determined by the header bidding line item could be *higher* than the bid itself. Indeed, publishers had an incentive to set floor prices that way, and there is evidence that publishers did just that.[699] Below, I correct the Plaintiffs' analysis by asking and answering three questions.

368. *First*, does the AdX second-price auction with a floor price based on winning header bids inevitably advantage AdX bidders? *No*.

[698] That is, the highest willingness-to-pay among bidders on AdX and the non-Google exchange are equal.

[699] *See also* Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -506 ("We've anecdotally heard from some publishers that they inflate the value CPM of remnant line items."); Presentation, "First-price bidding" (Aug. 12, 2019), GOOG-DOJ-11406673, at -677 ("How boost works[:] The publisher inflates the HB bid before sending it as a floor to AdX[.] This is done to increase Adwords cost and to provide a better comparison between Adwords and header bidder bids[.]"). *See also* Email from [redacted] to [redacted] et al., "Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019" (Aug. 13, 2019), GOOG-DOJ-09713317, at -319 ("Today, these [publisher-]inflated CPMs are used to provide price pressure for AdX [...] In practice, [...] many publishers [...] apply a boost to Header Bidding bids"); Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -777 ("Last look [...] incentivizes pubs to inflate ('boost') the floor sent to AdX"); [redacted]

369. Header bidding auctions are typically first-price auctions. As I have argued in Section III.C.3.b, in first-price auctions, it is optimal for bidders to shade their bids: that is, an exchange that bids in a header bidding auction will bid less than the value of its highest bidder. For example, it might be optimal for an exchange to bid $1 for an impression in a first-price header-bidding auction when its highest bidder's value is $1.50. *If* the publisher uses that $1 bid as its floor price in AdX's second-price auction, the header bidder would lose the auction to an AdX advertiser who truthfully reports a value of $1.10, even though the AdX advertiser has the lower value. *But* the publisher could (and would have incentive to) reverse this effect by choosing a higher floor price for the AdX auction, which it can do by triggering a header bidding line item with a value CPM inflated above the header bidding offer. For example, if it triggers a line item with value CPM equal to $1.50—so that the AdX floor price is $1.50—then the AdX advertiser has *no advantage* from the so-called "last look": it loses because its bid of $1.10 falls below the floor price. Indeed, with a floor price of $1.50, an AdX bidder can only win if it is willing to pay at least $1.50 for the impression, which is the same as the bid it would need to win a (hypothetical) second-price auction combining bids from header bidders and AdX. If the publisher chose an even higher floor price than $1.50, the AdX bidder would be *disadvantaged*. As Professor Weinberg notes, publishers could readily configure header bidding to more generally implement this bid adjustment procedure.[700]

[700] A publisher could set the value CPM of each header bidding line item as a function of the header bid or by passing a header bid inflated directly on the end user's browser. *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 141 ("Publishers had the ability to increase the clearing price passed on from their header bidding setup to DFP, such as with a multiplier, or an added value."). *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶¶ 14 ("Up until at least December 2021, the winning bid from the Header Bidding auction was typically used to trigger a specific line item that the publisher had booked within Google's ad server (most commonly a remnant line item), and […] the Value CPM of that line item could represent the winning Header Bidding bid as a floor in the AdX auction (prior to September 2019) or as a competing bid in the Unified First Price Auction (from September 2019 onwards)."), 11 ("Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other

370. *Second,* does a publisher using header bidding have an *incentive* to increase AdX floor prices as just described, reducing or reversing the claimed "last look" advantage? *Yes.*

371. In a second-price auction, it is always profitable for a publisher to trigger an AdX floor price higher than the highest header bid. In the above example, the publisher who has a $1 winning header bid in hand is *guaranteed* to sell the impression for $1 simply by allocating the impression to the winning header bidder. *If* the publisher uses that $1 bid as its floor price in AdX's second-price auction, then the publisher gains only when at least *two* bids in AdX's second-price auction are above $1. *But* the publisher can always do better: for instance, by triggering a header bidding line item with a value CPM of $1.01, the publisher gains when at least *one* bid in AdX's second-price auction is above $1. Of course, the publisher can do even better by setting the *optimal* value CPM based on the publisher's probabilistic assessments about bids in the AdX auction or based on experiments that it may run. A textbook result in auction theory implies this: it is always optimal for a seller to set an auction's floor price higher than its value for the item outside the auction.[701]

372. *Third*, if a publisher chooses floor prices *optimally* to maximize its average revenues and advertisers bid to maximize their profits, are AdX bidders always advantaged compared to header bidders as a consequence of the so-called "last look"? *No.*

reasons."). *See also* Asmaâ Bentahar, "Bid Adjustments Simplified: Run Fair Auctions with no Hassle," Pubstack (May 2, 2021), https://www.pubstack.io/topics/bid-adjustments-simplified ("Alternatively, this next one will slightly increase all returned CPMs, giving Prebid an edge in the competition against GAM").

[701] Krishna, V. (2010). *Auction theory* (2nd ed.). Academic Press, at 23 ("[A] revenue maximizing seller should always set a reserve price that exceeds his or her value.") (emphasis omitted).

373. In Theorem 3 below, I show that, in a standard model of auctions and the same one that Plaintiffs' expert Professor Weinberg recommends,[702] when all participants maximize their payoffs, AdX bidders enjoy no advantage over header bidders arising from the so-called "last look."

374. **Theorem 3:** Suppose that (i) a publisher sells an impression to a fixed set of bidders, including bidders on AdX; (ii) the publisher does not know each bidder's value for the impression; (iii) bidders do not know each other's values for the impression; and (iv) all participants understand the rules and have the following information:

a. Each bidder knows its own value for the impression.

b. Each bidder's value is drawn from a commonly-known probability distribution[703] and is statistically independent from other bidders' values.[704]

[702] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 40 ("In the context of the auctions for online display ads, I believe it is best to conduct the analysis assuming that advertisers have independent private values for impressions. The only exception is when I consider the potential impact of Enhanced Dynamic Allocation on direct deals in Section IV."). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions (the only exception is when identifying a potential impact of Enhanced Dynamic Allocation on direct deals in Section IV). This is a simplifying assumption (it is likely that no auction in real life purely abides by the independent private values model) that makes the analysis more tractable, and it is a sensible assumption to make since (a) internal Google documents demonstrate that Google assumes this as well (e.g., GOOG-AT-MDL-004016180), (b) bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's) and since the bidders do not know each other's identities, even if they learn about others' bids it could possibly not be that useful towards determining their own valuation."). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), Appendix D at ¶ 6 ("Consider an auction in the independent private values model where there are 2 bidders whose values are each drawn independently and uniformly from [0,10].").

[703] I assume that this probability distribution satisfies a technical condition called "Myersonian regularity," described in Section XV.D, which ensures that the optimal mechanism for selling the impression is an auction.

[704] This implies that each bidder and the publisher can make a probabilistic assessment about other bidders' values and estimates of other bidders' values would not be changed upon learning one bidder's value.

c. Each bidder determines a bid as a function of its value to maximize its surplus from the impression, given its probabilistic assessments about the bids of other bidders.

Then, if the publisher chooses revenue-maximizing floor prices, it earns exactly the *same* expected revenue with a single unified second-price auction with all bidders participating as it would with first-price header bidding auction followed by a second-price auction for AdX.[705] Moreover, given those optimal floor prices, each header bidder has the same chance of winning and the same expected surplus as an identical bidder on AdX.

375. A proof of Theorem 3 can be found in Section XV.D. The intuition for the result is that, after receiving the bids from header bidders, publishers can ensure that a bidder on AdX wins only if its value is higher than those of all header bidders, by committing to choose a floor price in the second-price auction that reflects the highest value among header bidders. This means the publisher can ensure that the highest value bidder wins under each auction format, resulting in the efficient allocation above the floor price and—by the Revenue Equivalence Theorem that applies in this setting (see Paragraph 254)—the same profit-maximizing revenues under the two auction formats.

376. Theorem 3 demonstrates that the so-called "last look" need not necessarily offer any advantage to AdX (or any other ad exchange) when publishers boost header bids and all parties respond to their incentives. Professor Weinberg acknowledges publisher incentives to boost header bids, but argues that such behavior requires a degree of

[705] More specifically, the theorem considers an auction game in which the publisher moves first, the header bidders second, and AdX bidders last. The publisher selects header bidding floor prices and a function to map each header bid it receives into a floor price for the AdX auction; the header bidders select their bids; and then finally the AdX auction is run.

"sophistication."[706] However, because ad revenue is a significant source of income for many publishers, I expect them to reason through and experiment with various sensible pricing strategies. As I previously noted, documents suggest that publishers *did* "inflate[] the HB bid before sending it as a floor to AdX"[707] in practice and, as Professor Weinberg writes himself, some of the most popular header bidding tools readily provided publishers with the functionality to boost bids.[708] Although there could have been some publishers

[706] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 165 ("The impact of Dynamic Allocation with sophisticated publishers who cleverly set Value CPMs is less clear-cut. On one hand, if sophisticated publishers only slightly inflate the Value CPM of the winning header bid, then the above conclusions continue to hold for exactly the same reasons. On the other hand, if sophisticated publishers significantly inflate the Value CPM of the winning header bid due to Dynamic Allocation and would not have set such an inflated reserve on AdX in absence of Dynamic Allocation, then the cost of this inflated reserve might outweigh the benefits highlighted above."). *See also* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 221 ("This is my opinion in aggregate, after considering that some publishers used default options while others were sophisticated and increased Value CPMs of header bids to boost AdX's reserve.").

[707] *See also* Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -506 ("We've anecdotally heard from some publishers that they inflate the value CPM of remnant line items."); Presentation, "First-price bidding" (Aug. 12, 2019), GOOG-DOJ-11406673, at -677 ("How boost works[:] The publisher inflates the [header bidding] bid before sending it as a floor to AdX[.] This is done to increase Adwords cost and to provide a better comparison between Adwords and header bidder bids[.]"). *See also* Email from R. Srinivasan to B. Bender et al., "Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019" (Aug. 13, 2019), GOOG-DOJ-09713317, at -319 ("Today, these [publisher-]inflated CPMs are used to provide price pressure for AdX [...] In practice, [...] many publishers [...] apply a boost to Header Bidding bids"); Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -777 ("Last look [...] incentivizes pubs to inflate ('boost') the floor sent to AdX"); [redacted]

[708] A publisher could set the value CPM of each header bidding line item as a function of the header bid or by passing a header bid inflated directly on the end user's browser. *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 141 ("Publishers had the ability to increase the clearing price passed on from their header bidding setup to DFP, such as with a multiplier, or an added value."). *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶¶ 14 ("Up until at least December 2021, the winning bid from the Header Bidding auction was typically used to trigger a specific line item that the publisher had booked within Google's ad server (most commonly a remnant line item), and […] the Value CPM of that line item could represent the winning Header Bidding bid as a floor in the AdX auction (prior to September 2019) or as a competing bid in the Unified First Price Auction (from September 2019 onwards)."), 11 ("Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons."). *See also* Asmaâ Bentahar, "Bid Adjustments Simplified: Run Fair Auctions with no Hassle," Pubstack (May 2, 2021), https://www.pubstack.io/topics/bid-adjustments-simplified ("Alternatively, this next one will slightly increase all returned CPMs, giving Prebid an edge in the competition against GAM").

who did not boost header bids in this manner, it is possible that boosting would have small effects on those publishers' revenues, or that they addressed the differences between exchanges' auction formats through other means, such as increasing the AdX specific floor price or using rebates (which I discuss more generally in Section XIV.E.3). Furthermore, Plaintiffs' analyses provide no evidence to support their assumption that publishers did not systematically boost header bids.

377. Even when the assumptions of Theorem 3 are not satisfied, publishers can use similar strategies to inflate value CPMs into AdX if they wish, or other strategies that are even more profitable for them. Any residual advantage or disadvantage for AdX then arises as a result of publishers' decisions to maximize profits, rather than from the so-called "last look."

3. Plaintiffs' and Their Experts Significantly Overstate the Effects of Line Item Caps

378. Since 2013, Google's ad server has limited the maximum number of active line items to 61,000[709] in order "to protect the health of the product [and] the performance of [Google's] system."[710] Plaintiffs and their experts assert that this limit "throttles

[709] Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -198 ("Number of active line item limits (please note: Each creative-level targeting criterion counts also toward the active line item limit[:] 61,000").

[710] Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -195 ("In 2014, we have started enforcing limits on the creation of certain entities in the DFP Ul, e.g. Lls per order, Creatives per LI, Ad units for placements and Targeting attributes in line items. [...] Limits are necessary in order to protect the health of the product, the performance of our system, and are ultimately for the benefit of all publishers and the performance of their Ul (unlimited number of entities can affect the Ul in terms of load time and serving in some cases). NOTE: Our limits are much higher than those that existed previously in DART.").

publishers' use of header bidding by artificially capping publishers [*sic*] use of 'line items'[...]" but their conclusions are flawed for several reasons.[711]

379. *First,* in dismissing Google's rationale for the caps as "pretextual," Professor Gans mischaracterizes the documents on which he claims to rely.[712]

a. Professor Gans argues that Google's "true motives" could not have been grounded in technical constraints because, in a 2017 email exchange, one Google employee remarked that "we may need to figure out the real hard limit."[713] However, the entire email exchange seems to concern a "system [] instability" and infrastructure issue arising from a publisher (CBS) that had over *one million* line

[711] Fourth Amended Complaint ¶ 389. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 634 ("The third way in which Google impaired the use of its ad server products was by imposing restrictions on 'line items' to limit the use of Header Bidding by publishers. Line items were ad server settings that publishers used in order to customize their ad servers, most notably to enable Header Bidding. As explained previously, Header Bidding was a key facilitator of competition for the inventory between AdX and third-party exchanges. Google imposed restrictions on the number of line items publishers could use ('line item caps'), in an effort to limit the adoption of Header Bidding."). *See also* Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 149 ("Google also could use the ad server technology to impede the use of Header Bidding through line-item caps.").

[712] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 646 ("Google's motive in imposing limits to the number of line items available to publishers, and denying requests for those limits to be raised, was to restrict competition in the ad exchange market by making Header Bidding more difficult and costly to the largest and most important publishers. While Google offered various technical explanations for the caps, these were pretextual.").

[713] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 647 ("Google communications make clear that the true motives for limiting line items was to prevent Header Bidding, rather than technical costs. A 2017 email explains that Google needs to "figure out the real hard limit.").

items at the time.[714] While Professor Gans interprets the employee's comment as refuting the need for a cap, the full conversation confirms in detail how [redacted] configuration presented "a real danger for the stability of [Google's] system" that required setting guard-rails to ensure consistently stable performance.[715]

b. Professor Gans then interprets another email exchange to suggest that reducing the number of line items would "limit Header bidding."[716] But that same exchange notes that "10k is more than enough line items for a HB setup, so I don't think ALI limits will be a motivator to move pubs off," and finds that reducing the number of active line items would likely have little impact on publishers, with

[714] Email from [redacted] to [redacted], "Re: [URGENT]: [redacted] hit their total line item limit" (March 29, 2017), GOOG-DOJ-15442474, at -474 to -479 ("[redacted] wants to increase their total LI limit from 1,000,000 to 1,200,000 asap [...] 1. Why do inactive/archived lines count towards the total limit? Infrastructure. Is what it is for now, unfortunately [...] Confirmed that header bidding caused this: [...] That order has 300k lines! [...] That sample order I found has around 300-400 lines but the story is that there are a bunch of orders related to header bidding causing this. [...] I don't think we should subject our system to instability due to this type of setup. [...] we may even need to get eng/pm director-level input to figure out how to handle and communicate this to pubs. [...] Non-active line items don't stress ad-serving or IM/F but they would stress Frontend, API and reporting. For Frontend, we have to search, sort and filter across all of these line items. So this does affect performance. External API requests would also need more resources and time to process. Imagine an external API script dumping all line items (which pubs often do). [...] Reporting also needs to deal with any line item that ever served an impression so I am assuming more line items would mean greater load but don't know specifics. [...] [redacted] is right in that this represents a real danger for the stability of our system.").

[715] Email from [redacted] to [redacted], "Re: [URGENT]: [redacted] hit their total line item limit" (March 29, 2017), GOOG-DOJ-15442474 ("[redacted] is right in that this represents a real danger for the stability of our system.").

[716] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 649 ("In 2018, Google employees again explained their methods to limit Header Bidding. In a 2018 email exchange, it seems that the costs for Google of doing so are limited ('(some) cost on us') and that resources can be saved some other ways ('purging active LI that aren't really active is something we should do anyway').").

about a quarter of line items that were active for six months or longer having no impressions in six months.[717]

c. Professor Gans also claims that "Google declined many publishers' access to additional line items despite its ability to do so,"[718] but he forms his conclusion on anecdotal evidence that ignores the many exemptions that Google *did* provide to publishers. In the data I have reviewed, at least [redacted] publishers are approved to set more than 61,000 active line items. For example, both [redacted] and [redacted] can set up to [redacted] active line items.[719]

380. Other documents cited by Professor Gans confirm that Google had legitimate concerns about how publisher configurations imposed costs and created a strain on Google's infrastructure. In an email thread he cites, Google employees note that "we h[a]ve seen growth in inventory complexity (steadily) for >2 years. This required us to refactor our

[717] Email from N. Korula to [redacted] et al.,, "Re: ALI limits/ fee" (Oct. 8, 2018), GOOG-DOJ-15127000 from -000 to -004 [redacted]

[718] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 651 ("Google only partially provided publishers with more line items via exceptions. In August 2018, Google launched a new process for publishers to request line item limit exceptions. Google notes that requesting an exception does not guarantee a publisher access to more line items. Google declined many publishers' access to additional line items despite its ability to do so[...].").

[719] These findings follow from my analysis of the "GAM Publisher Active Line Items Data" at GOOG-AT-EDTX-DATA-001116098, in which I logged the number of NetworkIds with a Limit larger than 61,000, in addition to the Limit for entries named "[redacted] The relevant code is in code/line_item_capping.py in my supporting materials, and the output is saved in code/logs/line_item_capping.txt.

serving infrastructure for retrieving LI/CLT. The new infrastructure scales well, all we need to support continued growth is more resources (which itself is an issue...)[.] [redacted] We mitigate and resolve these as they come up -- but it highlights the cost of allowing unrestrained growth. It comes with a cost, and we have to impose limits."[720] An additional document I reviewed noted a publisher with "[redacted] active [line item]s that put a huge strain on [Google's] system []."[721] Altogether, Professor Gans' assertion that technical considerations were "pretextual" is inconsistent with the themes of these emails and documents.

381. *Second,* despite Plaintiffs repeated emphasis of the line item limit, the data I analyze suggests that few publishers approach this threshold in practice. While Professor Pathak states that some "large publishers needed to create thousands of line items" in order "to run their Header Bidding arrangements," he ignores that GAM's 61,000 line item limit *did* accommodate "thousands" of line items.[722] In October 2016, when Google evaluated the effect of strictly enforcing the line item cap, it found that just four publishers would be affected.[723] My analysis of GAM data from April of 2024 is consistent with that document and found that over [redacted]% of publishers utilize less than [redacted]% of their active

[720] *See* Email from [redacted] to [redacted] et al., "Re: Ultraprio - Increase the ALI for Turner" (Sep. 26, 2018), GOOG-TEX-00090969, at -970.

[721] *See* Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -202.

[722] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 149 ("As I understand, to run their Header Bidding arrangements, some large publishers needed to create thousands of line items, which represent potential Header Bidding price points in DFP.").

[723] Comms Doc, "Limits Enforcement" (Feb. 16, 2018), GOOG-DOJ-09494195, at -201 ("Please note that as of Oct 2016, we only have [redacted] DFP publishers exceeding this limit and those are being closely monitored by Eng to ensure no negative repercussions are happening on our serving infrastructure.").

line item limit.[724] If the granularity of line items were as significant of an issue as Plaintiffs pose, one would expect the number of publishers near capacity to be substantially higher. Documentation of the popular header bidding wrapper Prebid.js echoes this finding, with the highest granularity option using just 2,001 line items per exchange.[725]

382. *Third,* as noted by Professor Gans, in 2021, Google began to roll out Header Bidding via Yield Groups (HBYG), which rendered it unnecessary for publishers to set up thousands of line items to call header bidding exchanges.[726] Professor Gans criticizes HBYG for "limited functionality" as it "only worked with the Prebid wrapper,"[727] but Prebid is "the most widely used header bidding 'container' or 'wrapper' on the web."[728] Moreover, publishers could use HBYG in conjunction with header bidding to reach the few partners who might not use Prebid. He also criticizes HBYG for being "confidential" during its

[724] These findings follow from my analysis of the "GAM Publisher Active Line Items Data" at GOOG-AT-EDTX-DATA-001116098. I calculate each publisher's percent utilization of its cap by dividing its "Active including creatives" by its "Limit," noting that over [redacted]% of values in this list are less than [redacted]%. The relevant code is in code/line_item_capping.py in my supporting materials, and the output is saved in code/logs/line_item_capping.txt.

[725] Prebid.org, "Prebid.js Ad Ops" (accessed Jul. 24, 2024), https://docs.prebid.org/adops/price-granularity.html#prebid-default-price-granularities.

[726] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 657 ("With HBM, publishers no longer had to set up hundreds of line items for Prebid and could instead install Header Bidding via Yield Groups, similar to Open Bidding Yield Groups.").

[727] *See* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 658 ("HBM did not fully solve the issue of publishers requesting line item limit exceptions as it had limited functionality. HBM only worked with the Prebid wrapper and GPT tags, and as of March 2023, not all Header Bidding Ad Networks were supported by HBM.").

[728] Prebid.org, "About," (accessed Jul. 25, 2024), https://prebid.org/about.

pilot,[729] even though it is not unusual for firms to keep programs in pilot confidential. In April of 2022, HBYG was announced as a public product.[730]

383. *Fourth,* Plaintiffs' and their experts' descriptions of how header bidding functions are factually incorrect. Plaintiffs pose a hypothetical in which "a publisher receives a header bidding exchange bid of $4.29, but only has a pre-existing line item with a price of $4.20, then the publisher's Google ad server rounds down the header bidding bid to the line item with the next closest price, e.g., to the line item with the price of $4.20,"[731] asserting that "[f]ewer line items cause publishers' bids from header bidding exchanges to be rounded down more often," leading publishers to "receive less revenue when a header bidding exchange wins."[732] This is wrong for two reasons.

i. Rounding a winning header bid does not affect the payment amount from the exchange. When a header bidding exchange offers $4.29 and wins, the publisher receives $4.29 regardless of whether the publisher triggered a $4.30 or $4.20 line item in GAM.[733, 734]

[729] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 658 ("Moreover, HBM adoption was slow as it was available in GAM 360 only and remained confidential as only a couple of publishers were progressively enrolled. Moreover, Google tried to keep HBM confidential but was noticed by some bidders.").

[730] *See* Google Ad Manager Blog, "Improved header bidding support in Google Ad Manager" (Apr. 27, 2022), https://blog.google/products/admanager/improved-header-bidding-support-in-google-ad-manager/.

[731] Fourth Amended Complaint ¶ 390. *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 661 ("If a publisher receives a bid of $5.20 from an exchange using Header Bidding, but the publisher only has a pre-existing line item with a price of $5, then Google's ad server rounds down the Header Bidding bid to the line item with the next closest price (in this case, $5).").

[732] Fourth Amended Complaint ¶ 394.

[733] *See, e.g.*, Prebid.org, "Price Granularity" (accessed Jul. 12, 2024), https://docs.prebid.org/adops/price-granularity.html ("Important: Rounding does not impact the price paid, only the auction on the ad server. For example, if your bid for 2.75 is rounded down to 2.50 and wins on the ad server at 2.50, you will be paid 2.75.").

[734] This example omits the additional issue of payment discrepancies in header bidding. As I discuss further in Section X, there were reports of payment discrepancies, where publishers' expected receipts from header bidders did

ii. Plaintiffs' allegations overlook publishers' incentives to increase the AdX floor price above header bids. In the hypothetical described above, *even if* the line item cap required the publisher to use less granular floor prices, Plaintiffs are wrong to assume that the publisher rounds down to $4.20. Instead, as I discussed above, the publisher would be incentivized to *round up* the header bid to $4.30 or to inflate the bid even more. Because header bidding was a publisher-configured technology, publishers could trigger line items with value CPMs higher than header bids. Additionally, header bidding wrappers such as Prebid.js provided publishers with the functionality to round bids *up* rather than *down*.[735] As I have explained in Section XIII.D.2, configuring header bidding to trigger increased floor prices is consistent with publishers' incentives and can make bids from header bidding exchanges *more* competitive, not "less competitive compared to the bids from Google [...]."[736]

4. Plaintiffs and Their Experts Misinterpret Google's Experiments on "Last Look"

384. Plaintiffs and their experts' claim that Google experiments predicted that removing the so-called "last look" would have reduced the revenues of AdX, Google Ads, and DV360.[737] Plaintiffs' and Professor Gans' allegations appear to refer to experiments

not match the eventual payments. *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24 ("Header bidding is also not transparent because, although the publisher 'accepts' the impression at a certain price, the header bidder may not actually pay the sum indicated in its bid.").

[735] Prebid.org, "Prebid.js Publisher API Reference" (accessed Jul. 24, 2024), https://docs.prebid.org/dev-docs/publisher-api-reference/setConfig.html#setConfig-Cpm-Rounding ("Prebid also allows setting a custom rounding function. This function will be used by Prebid to determine what increment a bid will round to.").

[736] Fourth Amended Complaint ¶ 394.

[737] *See* Fourth Amended Complaint ¶ 382 ("Truly giving up Last Look would have cost Google too much; Google predicted a 10 percent hit to DV360's revenue and at least a 30 percent decrease in Google Ads' revenue."). *See also* Expert Report of J. Gans (Jun. 7, 2024), at ¶ 609 ("Google conducted experiments to assess the significance of the

summarized in an internal presentation during the transition to the Unified First Price Auction.[738] However, that experiment tested just one way in which the so-called "last look" could be removed, which was to remove remnant line items from the calculation of the floor price for the AdX second-price auction.[739] By itself, that change would reduce AdX's win rate: it would reduce AdX floor prices (because remnant line items could no longer increase those floor prices) and thus reduce the AdX clearing price (because the AdX floor sometimes set the clearing price in the AdX second-price auction). And because, in the experiment, the AdX clearing price was compared to other exchanges' bids, lower AdX clearing prices would translate into fewer wins for AdX bidders.

385. Plaintiffs treat the experimental results as evidence of how the industry would have evolved without the so-called "last look," but this analysis incorrectly assumes that the only way for Google to "remove" the so-called "last look" was the one tested in the experiment.

'Last Look' advantage before migrating to a first-price auction system. An internal document shows an [redacted]% decrease in AdX revenue and [redacted]% decrease in impressions due to giving up 'Last Look.'").

[738] Professor Gans cites an email exchange dated August 29, 2019 (*See* Email from [redacted] to [redacted] et al., "Re: Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019," GOOG-TEX-00682264). The relevant quote states "On second-price traffic, this results in a [redacted]% decrease in AdX revenue [...] and a [redacted]% decrease in impressions." Although it is not possible to be certain, that email appears to be reporting results from the same analyses that were presented in a slide deck dated September 3, 2019 and titled "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-AT-02204351. Slide -382 of the presentation, titled "Last look removal (AdManager app & web)," contains near-identical text: "On second-price traffic, this results in a [redacted] decrease in AdX revenue and impressions."

[739] The title of the experiment cited in Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-AT-02204351, at -359 is [redacted] which is the basis for my understanding that remnant line items were removed from the calculation of the AdX clearing price.

386. The same document contains evidence that there were alternatives: AdX buyers could win more often by using a "Bid Translation service" to optimize their bids to compete against bids from other exchanges in a subsequent first-price auction.[740] And when Google ultimately removed the so-called "last look" over remnant line items, it did choose an alternative: it transitioned to the Unified First Price Auction, which avoided the adverse impacts of the change tested in the experiment.[741] As Plaintiffs concede themselves, after AdX *optimized* its bids to better compete against other exchanges, the changes were "revenue neutral" compared to the period in which it had a "last look,"[742] which suggests that there was no inherent advantage to "last look."

387. Moreover, as observed by the employee in the email exchange that Professor Gans cites, there are a number of short-term effects that make it "difficult" to interpret the results of the experiment. The employee notes various challenges, such as there being "[s]ome buyers [that] are experimenting with 'smart' / first-price bidding algorithms," and publishers who "inflate the CPMs of Header bidding line items."[743] These are typical limitations of short-run experiments. After participants adjust to the new incentives of the

[740] *See* Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-AT-02204351, at -362.

[741] I discuss Google's transition to the Unified First Price Auction further in Section XIV.

[742] *See* Fourth Amended Complaint ¶¶ 379-82 ("Several years later, in 2019, Google publicly announced that exchanges in Exchange Bidding would no longer be able to trade ahead of header bidding exchanges[.] […] Google compounded this Exchange Bidding advantage with a new secret bid optimization scheme that allowed Google to recapture the advantages it had under Last Look[.] […] Internal Google documents reveal that these changes were [redacted] for DV360 [redacted] percent) and Google Ads [redacted] percent).").

[743] Email from [redacted] to [redacted] et al., "Re: Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019" (Aug. 29, 2019), GOOG-TEX-00682264, at -265 ("Some buyers are experimenting with 'smart' / first-price bidding algorithms. Quantifying the effect of these is difficult, since buyers are using the transition period to develop and test new algorithms. Additionally, there is a long-term benefit of publishers having reduced incentives to inflate the CPMs of Header bidding line items. Today, these inflated CPMs are used to provide price pressure for AdX, which can result in increased publisher revenue. In a first-price auction, such inflation can only lead to reduced publisher revenue, so this inflation is expected to decrease over time.").

first-price auction, the actual revenue outcomes may differ from the outcomes in the experiment.

XI. RESERVE PRICE OPTIMIZATION: INCREASING PUBLISHER REVENUES IN "THIN" AUCTION MARKETS

A. Overview

388. Reserve Price Optimization (RPO) was an AdX feature designed to help publishers "earn the most money possible, with the least complexity."[744] Launched as Optimized Pricing in April 2015,[745] RPO increased the floor price sent to a bidder on AdX when Google predicted—based on historical bid data—that a higher floor price would increase publisher revenues.[746] By increasing publisher yields, RPO also incentivized publishers to make more inventory available for programmatic auctions.[747] Google updated its RPO features over time and found that each update increased publisher revenues.[748] Other

[744] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -154 ("Our goal has always been to help publishers thrive and create sustainable businesses with advertising: to earn the most money possible, with the least complexity, all while providing users the best experience. [...] These new features will help our publisher partners grow their revenue and give programmatic buyers greater access to premium inventory.").

[745] Email from [redacted] to [redacted], "Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers" (Nov. 12, 2015), GOOG-DOJ-07235914, at -916 ("In April, we launched a simple pricing model that sets prices based on inventory features such as web property and ad unit.").

[746] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 25, 2016), GOOG-DOJ-04937154, at -158 ("So in recent months, we've been working on optimized pricing technology that algorithmically sets floor prices in the Open Auction to increase publisher revenue. With optimized pricing, we use event level data available from previous auctions to predict what the bids will be on certain queries, and adjust the floor price accordingly on behalf of the publisher, subject to their settings.").

[747] Presentation, "Reserve Price Optimization, Optimized Private Auctions & Dynamic Revenue Sharing" (Mar. 23, 2018), GOOG-AT-MDL-004242638, at -639 ("When publishers' programmatic yield grows, they make more inventory available to buyers.").

[748] Presentation, "AdX Dynamic Price" (Dec. 11, 2014), GOOG-DOJ-13199910, at -930 ("AdX pubs [...] Total [redacted]% (with GDN opt out)"; Launch Doc, "AdX Dynamic Price Optimization V2" (Sep. 16, 2015), GOOG-DOJ-13209957, at -959 ("The cookie model seems to add about [redacted]% incremental revenue."); Email from [redacted] to [redacted], "Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers" (Nov. 12, 2015), GOOG-DOJ-07235914, at -915 ("Between April and October we launched and improved new systems to dynamically set auction reserve prices for AdX sellers. After three launches and a three month study of the impact on buyers and sellers, the Reserve Price Optimization (RPO) program now generates a total annual revenue lift of [redacted], or [redacted]% of total network (AdX+AdSense publisher) revenue (rasta)!"); Launch Doc, "Online Reserve Price Optimization" (Sep. 14, 2017), GOOG-DOJ-13211589, at -590 ("The launch candidate increases network-wide revenue by +1.11%").

sell-side intermediaries (including [redacted] Magnite (formerly Rubicon), [redacted], and [redacted]) implemented features similar to RPO.[749]

389. Plaintiffs allege that RPO misled advertisers because it "meant that the auction did not operate as a sealed second price auction as Google had advertised."[750] But, as noted by Plaintiffs' experts Professors Weinberg and Pathak, a second-price auction *remains* a second-price auction with or without a program like RPO in place.[751] RPO was a service to automate and improve the ordinary publisher task of setting floor prices in a second-price auction; it did not change the format of the AdX auction.

390. Plaintiffs and their experts allege that some publishers and advertisers were harmed by Google's alleged "concealment" of RPO.[752] But RPO was designed to increase

---

[749] [redacted] Rubicon Project, "Maintaining the Equilibrium: How Dynamic Price Floors Preserve the Integrity of the Automated Advertising Ecosystem" (2014), GOOG-AT-MDL-011234683, at -696 ("Rubicon Project gives sellers the opportunity to employ a sophisticated Dynamic Price Floor algorithm which works above a seller's hard floors."); [redacted]

[750] Fourth Amended Complaint ¶ 532.

[751] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 285 ("Note that each individual auction in isolation is still a second-price auction with reserve [...] .") (describing an example of the operation of RPO); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 191 ("As a result, even though AdX ran a second-price auction at the time [...] .") (discussing RPO's effects on advertiser bidding behavior).

[752] *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 272 ("The negative effects of RPO to advertiser payoff, and possibly some publishers' revenues, at least partially stem from Google's concealment of the conduct."); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 191 ("As a result, by concealing RPO, Google interfered with the publishers' revenue maximization [...] As a result, both the publishers and the advertisers could act accordingly if they knew RPO was in effect. This potentially reduces their revenue or payoff compared to what they could have been."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 360 ("It is my opinion that those undisclosed Google rule changes were contrary to the expectations of the auction participants and made it impossible for auction participants and competing exchanges to understand the rules that governed and applied to auctions run by Google, skewing decision-making and outcomes.") (discussing RPO, DRS v1, and Bernanke); Expert Report of J. Andrien at ¶¶

publishers' revenues, and it was successful at doing just that, with the initial versions of the program increasing publisher revenues on the order of [redacted]% and later updates increasing publisher revenues even further.[753] Plaintiffs and their experts do not provide actual examples of publishers claiming to be harmed by Google's RPO program, and nor am I aware of any such evidence. For advertisers, knowledge of RPO was not necessary for buy-side tools to optimize bids: in the absence of RPO, publishers could (and did[754]) use historical data to set floor prices,[755] so that a surplus-maximizing bidder would account for that possibility *both* when RPO was in place and before. Moreover, Google

34-35 ("For example, by concealing RPO, Google prevented publishers from effectively optimizing revenue, [...] Google also impacted advertiser behavior through its second-price auction representation and concealment of RPO.").

[753] [redacted]

[754] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155 ("Optimized pricing in the Open Auction automates the post-auction analysis and floor price updates that publishers are already doing and takes it a step further."); "Buyer incentives and reserve price optimization" (May 14, 2012), GOOG-DOJ-15588979, at -979 ("No expectation of bid privacy: other platforms like yield managers [] already give publishers full buyer bid information and therefore the buyers might not care").

[755] Indeed, as I discuss in Paragraph 397, in 2011, Google provided a Minimum CPM Recommendation Feature that helped publishers select such floor prices more effectively.

disclosed RPO roughly a year after it launched,[756] so any potential impacts from the alleged "concealment" would have been short-lived.

B. Importance of Floor Prices for Revenue Optimization

391. As I discussed in Section III.C.3.d, floor prices can help publishers increase their average auction revenues.[757] A publisher contemplating an increase in its floor prices faces a tradeoff between *increasing* the average price of each impression sold and *reducing* the probability of a sale because no bidders meet the price floor. Because a publisher typically does not know each advertiser's willingness to pay for an impression or even which advertisers are participating in each auction, publishers can use historical data on the sales of similar items or experimentation (or both) to set floor prices. For example, with historical data on bids received in a second-price auction, a seller can simulate the outcomes of auctions with alternative floor prices and choose the one that those simulations indicate would generate the highest revenue.[758] Alternatively, a publisher could run an experiment or "A/B test" of different floor prices on live auctions, and choose the floor price that led to the highest revenue in those experiments.[759]

[756] Jonathan Bellack, "Smarter optimizations to support a healthier programmatic market," Google Ad Manager (May 12, 2016), https://blog.google/products/admanager/smarter-optimizations-to-suppor/ ("In our experiments to date, we have applied optimized pricing to about [redacted]% of transactions, creating over [redacted]% lift in revenue for publishers using the Open Auction. As we expand our experiments with optimized pricing, we will monitor its performance to ensure advertisers continue to get great ROI."). *See also* Email from [redacted] to AdX Buy-Side Global Sales et al., "[ANNOUNCED] Smarter optimizations for DoubleClick Ad Exchange" (May 12, 2016), GOOG-DOJ-04934481, at -481 ("On Thursday afternoon we announced new optimizations for DoubleClick Ad Exchange, Optimized Private Auctions and optimized pricing in the Open Auction (aka RPO).").

[757] Note, however, that badly-chosen floor prices (*i.e.*, floor prices chosen too high) can *reduce* average auction revenues.

[758] *See* Ostrovsky, M., & Schwarz, M. (2023). Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy*, *131*(12), 3352-76..

[759] *See, e.g.*, Rhuggenaath, J., Akcay, A., Zhang, Y., & Kaymak, U. (2022). Setting reserve prices in second-price auctions with unobserved bids. *INFORMS Journal on Computing*, *34*(6), 2950-67.

392. Setting floor prices can be especially important in **thin** auctions, which are auctions with few competitive bidders.[760] This effect is seen most easily in second-price auctions because, in those types of auctions, a floor price can be effective at increasing the price of an impression *if and only if* there is a *single* bid above the floor price.[761] This is more likely to occur if there are fewer bidders, which is why floor prices are more important in thinner auctions.[762] Conversely, in second-price auctions with more bidders, floor prices are often less important because the price is more often set by the second-highest bid, rather than the floor price.

393. Although there are typically many advertisers participating in online display advertising auctions, there are several possible sources of auction thinness in display advertising.

394. *First*, not all bids from advertisers end up competing in the final auction for an impression. One reason for this is budget throttling, discussed in Section VI of this report, in which a buy-side tool selects a subset of its eligible advertisers for participation in each auction. Another reason, discussed in Section III.C.4 of this report, is that some buy-side tools submitted only one bid for an impression on behalf of multiple advertisers. Regardless of whether those bidding strategies benefit advertisers, they reduce the overall number of bidders participating for each impression, making it more likely that the publisher's floor price determines the clearing price of the impression.

760 In this context, "competitive" bidders are those bidders with values for items that are close to the average sale price for the impression.

761 The same result is true of first-price auctions, but for a different reason. In a first-price auction, each bidder's optimal bid depends on the number of competitive bidders in the auction and the floor price, with the optimal bid *lower* with *fewer* competitive bidders and *higher* when the floor price is *higher*. This makes the floor price a more important lever in first-price auctions with few competitive bidders.

762 *See, e.g.*, Reiley, D. H. (2006). Field experiments on the effects of reserve prices in auctions: More magic on the internet. *The RAND Journal of Economics*, *37*(1), 195-211 ("The gains to setting an optimal reserve price become very small as [the number of bidders] increases.").

395. *Second*, auction thinness can arise as a result of finer targeting, a possibility I first discussed in an academic paper in 2010.[763] Finer targeting allows an advertiser to more accurately identify the end users it seeks to reach with its advertising campaign. As a result, the advertiser bids *more* for impressions that meet its targeting criteria but participates in *fewer* auctions overall. Because the price an advertiser pays in an auction depends on the competitive environment (including the total number of advertisers competing for the impression), finer targeting can reduce overall auction revenues, even though the winning advertiser's bid is increased. For example, suppose that there are ten Ford dealerships running online display ad campaigns in Dallas–Fort Worth but only one operating in Plano. If an impression opportunity can be associated with someone shopping for a car in Dallas-Fort Worth, all ten might bid for the impression. But if each advertiser can identify that the end user lives in Plano, only one might be interested in bidding. As a result, the publisher's revenue from the impression can be lower in the second case than the first. This possibility had been discussed internally at Google, where some engineers called it the "Pricing Paradox."[764] Google engineers identified a "solution" to the "Pricing Paradox": more accurate floor prices.[765]

[763] Levin, J., & Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review: Papers & Proceedings*, *100*(2), 603-07, at 606 ("A second hazard of targeting is that it leads to thinner markets, which can create problems for accurate pricing.").

[764] Presentation, "The Pricing Paradox and solutions" (May 15, 2012), GOOG-DOJ-03366173, at -179 ("Pricing paradox[:] Targeting may create more value, less revenue").

[765] Presentation, "The Pricing Paradox and solutions" (May 15, 2012), GOOG-DOJ-03366173, at -188 ("Solution #2: Auctions with reserve prices [...] Compute revenue-maximizing reserve prices for each query bundle").

C. Google's Implementation of RPO

396. Over time, Google introduced a number of features to help publishers set floor prices leading to higher revenues in the AdX auction. In addition to improving programmatic revenues for its publisher customers, Google's revenue share model on AdX means that it also had an interest in setting floor prices leading to higher auction revenues.

397. In 2011, Google introduced a Minimum CPM Recommendation feature (internally called AdX Seller Reserve Price Optimization) for publishers.[766] The Minimum CPM Recommendation feature calculated optimal floor prices for each ad slot based on the previous week's bids on that ad slot and provided those as recommended floor prices in the publisher's AdX user interface, along with a graph of expected revenue for each floor price that the publisher could choose.[767] The publisher still needed to set the floor price manually in the AdX user interface based on its recommendations.[768] Google's early

[766] "AdX Seller Reserve Price Optimization" (Nov. 18, 2012), GOOG-DOJ-12439154, at -154 ("AdX Seller Reserve Price Optimization [...] Status: launched on Q2 2011").

[767] "AdX Seller Reserve Price Optimization" (Nov. 18, 2012), GOOG-DOJ-12439154, at -156 ("The approach is[:] 1. Process logs [redacted] moving window, updated daily) to extract first and second price statistics per ad slot[.] 2. Compute optimized floor prices for the next day"); Nemo Semret, "Introducing Minimum CPM Recommendations on DoubleClick Ad Exchange," DoubleClick Publisher Blog (Nov. 1, 2011), https://doubleclick-publishers.googleblog.com/2011/11/introducing-minimum-cpm-recommendations.html ("Today, we're happy to announce the launch of Minimum CPM Recommendations for DoubleClick Ad Exchange publishers. This feature automatically recommends an optimal minimum cpm for each eligible ad slot in the Ad Exchange auction. It also automatically generates a graph that provides better visibility into how different floor prices might affect a publisher's bottom line.").

[768] Design Doc, "Dynamic Floor Prices in AdX" (Aug. 20, 2012), GOOG-AT-MDL-010338120, at -120 ("Floor prices in AdX are set manually per ad unit (or, in the brave new adunitless world, per inventory rule). Either way, the min cpm is set by a human entering a number into a text box in the adseller UI.").

experiments found that publisher revenues increased approximately [redacted]% on the ad slots on which its recommended floor prices were adopted.[769]

398. Despite its Minimum CPM Recommendation feature and other strategies used by publishers to increase auction clearing prices,[770] by 2015, Google found that "even with all this effort, there [was] still a wide and persistent price gap between the bid and closing prices in the open auction across all [] publishers."[771] Some Google engineers referred to that gap as the "auction discount" because (in a second-price auction with optimal bids) it is the difference between the bidders' willingness-to-pay and the price it actually paid.[772] Google engineers measured the "auction discount" in 2015, and found that [redacted].[773]

[769] Nemo Semret, "Introducing Minimum CPM Recommendations on DoubleClick Ad Exchange," DoubleClick Publisher Blog (Nov. 1, 2011), https://doubleclick-publishers.googleblog.com/2011/11/introducing-minimum-cpm-recommendations.html ("Initial results with early beta testers indicate an average 20% revenue lift for adopted recommendations.").

[770] For example, one Google document notes that "publishers have created complex systems of publisher-set floors to close the gap. Unfortunately, these floors are hard to calculate manually, requiring ad ops teams to spend countless hours gathering data and running post-auction analysis to update pricing and priorities in their system. Some publishers have even resorted to more extreme methods like waterfalls between exchanges, which introduces latency that damages consumer experience and advertiser performance." Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155.

[771] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155.

[772] "Cookie based Dynamic Reserve Price Optimization (RPO) - mini PRD" (May 1, 2015), GOOG-DOJ-13203511, at -511 ("There is often a [redacted] a gap between what buyers are bidding and what they end up paying when the auction closes.").

[773] "Cookie based Dynamic Reserve Price Optimization (RPO) - mini PRD" (May 1, 2015), GOOG-DOJ-13203511, at -511 ("On average buyers pay less than half of what they bid. [redacted].").

399. In April 2015, Google introduced its AdX Dynamic Price feature (later simply called RPO, which is the term I will use for the program).[774] RPO *automatically* increased floor prices for publishers with the simple goal of helping publishers "earn the most money possible, with the least complexity."[775] Google engineers noted that RPO could also "encourage publishers to make more inventory accessible to the open auction."[776] RPO also responded to auction "thinness" created by the one-bid policies chosen by some buy-side tools participating in the AdX auction (see Section III.C.4). As one Google engineer noted, "A dynamic RPO floor effectively recovers some of the money lost by the fact that some bidders[,] say [C]riteo, aren't sending a second bid, they are just sending [one]."[777]

400. The first RPO model (sometimes called Per-Buyer RPO) increased floor prices on a per-buyer basis, using data on the buyer's bids on a publisher's ad slot from the previous day to estimate the floor price that maximized the expected auction revenue.[778] In October 2015, Google added Cookie-Based RPO, which used the previous day's data on a buyer's

---

[774] Email from [redacted] to [redacted], "Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers" (Nov. 12, 2015), GOOG-DOJ-07235914, at -916 ("In April, we launched a simple pricing model that sets prices based on inventory features such as web property and ad unit.").

[775] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -154 ("Our goal has always been to help publishers thrive and create sustainable businesses with advertising: to earn the most money possible, with the least complexity, all while providing users the best experience,"), -155 ("Optimized pricing effectively reduces the gap between the first price and closing price increasing publisher yield. Optimized pricing could encourage publishers to make more inventory accessible to the open auction as well as reduce complex setups with varying ad server priority and floor prices[.]").

[776] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155.

[777] "Notes: MTV/LON/CAM Strategy Summit" (Aug. 12, 2015), GOOG-DOJ-10572595, at -603.

[778] Email from [redacted] to [redacted], "[Launch 124105] Per-buyer dynamic reserve price optimization on AdX - full launch (+$[redacted] annual revenue from RTB buyers)" (Nov. 13, 2014), GOOG-DOJ-15419945, at -945; Presentation, "AdX Dynamic Price" (Dec. 11, 2014), GOOG-DOJ-13199910, at -920 ("Daily pipeline to compute pricing file based on 'yesterday' data), -925 ("Compute bid distributions for 'yesterday'[.] Find 'optimal' reserve prices[...]").

bids for a given cookie to estimate revenue-maximizing floor prices for impression opportunities associated with that cookie.[779] Cookie information was included in the RPO model because Google found that "[b]uyers often buy based on cookies, and cookies are an important part of the valuation of every query."[780] Per-Buyer RPO and Cookie-Based RPO could apply simultaneously to a single impression, in which case the higher of the two estimated floor prices would apply.[781] Overall, an RPO floor applied to fewer than [redacted]% of impressions.[782]

401. RPO only ever increased floor prices. It never set floor prices below the publisher's chosen floor price, even if Google estimated that floor price to be higher than optimal.[783] RPO increased floor prices to the level that maximized the publisher's predicted profits, subject to a constraint that the publisher's predicted match rate did not decrease by more than a certain amount.[784, 785] RPO automatically applied to publisher's Open Auction

---

[779] Email from [redacted] to [redacted], "[Launch 132012] Cookie based Dynamic Reserve Price Optimization on AdX - full launch" (Sep. 29, 2015), GOOG-DOJ-15423462, at -462 ("Launch Date 2015-10-05").

[780] "Cookie based Dynamic Reserve Price Optimization (RPO) - mini PRD" (May 1, 2015), GOOG-DOJ-13203511, at -511.

[781] Presentation, "AdX Dynamic Price V2" (May 26, 2015), GOOG-DOJ-14000011, at -012 ("Effective reserve is max of inventory-RPO, cookie-RPO price").

[782] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -160 ("A: Optimized pricing has no effect on most Open Auction queries - right now it affects the winning price for fewer than [redacted]% of impressions bought by RTB buyers."); "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273 ("Percentage of Impressions" table, "DYNAMIC_RESERVE" row, "Grand Total [...] [redacted]%").

[783] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155 ("Optimized pricing can only increase floors from where the publisher has them currently set. It will never lower a floor.").

[784] Email from [redacted] to [redacted], "Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers" (Nov. 12, 2015), GOOG-DOJ-07235914, at -916 ("Pick a reserve price that maximizes predicted revenue, constrained to limit the fraction of bids it eliminates (to preserve match rate).").

[785] At least initially, the match rate was constrained to dropping by no more than [redacted]%. *See* Design Doc, "Cookie Data in Reserve Price Optimization Pipeline" (Mar. 29, 2015), GOOG-DOJ-13200480, at -482 ("We simulate with the default values of the Dynamic Reserve Price Rule Computation parameters. These values are the ones used in

inventory, and there was no option for a publisher to opt out from RPO.[786] While Google had flagged as early as 2014 the possibility that a publisher's floor price might be modified on some impressions,[787] RPO was officially announced to the public on May 12, 2016, as Optimized Pricing.[788]

402. Initially, RPO exempted buyers that submitted at least two bids into the AdX auction.[789] Google engineers discussed this exemption internally as a way to "encourage buyers to declare two bids"[790] in the AdX auction, and it also addressed "concerns" about the interaction of Project Bernanke and RPO.[791] In November 2017, Google engineers noted that the exemption policy was "easy to circumvent by submitting a nominal (1 cent) minimum payment amount with every bid" and, as a result, changed the RPO exemption

the Inventory RPO that runs daily in production. [redacted])") (emphasis omitted).

786 Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155 ("The publisher does not get any new controls or and opt out and the feature works only in the background.").

787 "Ad Exchange auction model" (Aug. 24, 2014), GOOG-AT-MDL-C-000035250, at -250 ("The Google DoubleClick Ad Exchange may run limited experiments designed to optimize the auction. These experiments may include modifying the standard auction model or mechanics (e.g., a tiered, rather than second price auction); simulating ad calls and auctions; modifying the min CPM set by the publisher for an impression or otherwise adjusting publisher settings; or discounting certain bids submitted by buyers or otherwise modifying the priority of the bids submitted by buyers, in an effort to optimize the auction. Publisher's buyer/advertiser blocks will not be modified.").

788 Email from [redacted] to AdX Buy-Side Global Sales et al., "[ANNOUNCED] Smarter optimizations for DoubleClick Ad Exchange" (May 12, 2016), GOOG-DOJ-04934481, at -481 ("On Thursday afternoon we announced new optimizations for DoubleClick Ad Exchange, Optimized Private Auctions and optimized pricing in the Open Auction (aka RPO).").

789 "AdX Dynamic Price: Sell Side Review fact sheet" (Dec. 2014), GOOG-DOJ-13200158, at -158 ("AdX buyers buying on both AdSense and AdX publishers will be subject to dynamic reserve pricing, subject to a policy [...] to exempt buyers who submit two bids or second price themselves.").

790 "The case for encouraging buyers to declare two bids" (May 11, 2015), GOOG-DOJ-13201465, at -465 ("We can encourage buyers to declare 2 bids per auction for three reasons: 1) In the presence of an effective RPO (reserve price optimization for AdX), AdX is in position to give a discount in pricing a buyer who declares two effective bids and price itself. This strategy will directly incentivises buyers to declare their second highest bid.").

791 "AdX Dynamic Reserve Price: AFC Launch Review Follow up" (Dec. 1, 2014), GOOG-DOJ-13199603, at -603 ("This will address two major concerns: Interaction of Bernanke and dynamic pricing is eliminated, as GDN always submits two bids and thus is exempt.").

policy to exempt bidders that submitted second bids high enough to generate "a certain amount of revenue lift from self-pricing."[792] In particular, Google exempted buyers for which the calculated increase in revenue from the second bid it submitted (compared to submitting only its high bid) was [redacted] times greater than a simulated increase in revenue Google calculated would occur due to RPO (if it submitted only its high bid).[793] Google Ads was exempt from RPO under both policies.[794] Both RPO exemption policies are consistent with RPO being designed in part to respond to auction "thinness" created by buy-side tools' decisions to employ single-bid policies, which (as I discussed above) would otherwise lead to reductions in publisher revenues.

403. In May 2018, Google introduced another version of RPO called Online RPO.[795] Online RPO was motivated by "buyers' bidding behavior on high value re[]marketing cookies," where Google had observed that most impression opportunities for a given cookie arrived within a period of one hour, making it difficult to use the previous day's data to optimize

---

[792] Launch Doc, "RPO Exemption Policy V2 Launch Doc" (Nov. 14, 2017), GOOG-DOJ-13212948, at -948.

[793] Launch Doc, "RPO Exemption Policy V2 Launch Doc" (Nov. 14, 2017), GOOG-DOJ-13212948, at -948 ("We run a simulation pipeline where we remove min cpm payments from buyers and calculate the revenue extracted from each buyer for both the case where min cpm payments are on and off. After that we measure the performance over a week worth of simulations and whitelist for RPO exemptions only those buyers that yield a revenue lift of more than [redacted] times the lift of RPO (over both AdX and AdSense), or [redacted]%. [...] At serving time in dynamic price producer we set the RPO reserves to 0 if a buyer belongs to this whitelist.").

[794] "AdX Dynamic Reserve Price: AFC Launch Review Follow up" (Dec. 1, 2014), GOOG-DOJ-13199603, at -603 ("We are exempting bidders (adx 'buyer networks') who submit a second bid to the AdX auction from dynamic pricing which will effectively make all of GDN demand exempt from dynamic pricing."); "Reveal RPO floor to Adwords" (Nov. 11, 2019), GOOG-DOJ-14030931, at -931 ("Previously in second price AdX auctions, Adwords was exempt from RPO floor.").

[795] "2018 Sellside Launches Revenue Evaluation" (Jul. 19, 2019), GOOG-DOJ-13949282, at tab "Q2 2018," row 6 (noting launch date of 5.10.18).

floor prices.[796] Online RPO maintained a sliding window of the [redacted] most recent bids for each buyer-cookie pair in the "high bid range" (meaning that bid requests for the cookie had received bids higher than some threshold) and determined an optimal floor price based on bids in that window.[797] Online RPO operated at the same time as the other RPO models (*i.e.*, Per-Buyer RPO and Cookie-Based RPO), and the maximum of the RPO-computed reserves would apply.[798]

404. Google conducted several experiments to assess the impact of its various RPO models. Experiments from around the time of the launch of Per-Buyer RPO found that it increased Google revenue and publisher payouts on the order of [redacted]% compared to no RPO.[799] Later experiments found that the integration of Cookie-Based RPO increased

[796] Email from [redacted] to [redacted], "UPCOMING LAUNCH - Please review: [Launch 225406] Online Reserve Price Optimization" (Feb. 5, 2018), GOOG-DOJ-14421383, at -383 ("The key motivation of online RPO is to use buyers' bidding behavior on high value re-marketing cookies to set reserve prices. While analyzing bidding patterns on cookies, we observed that most queries for a given cookie arrive within a one hour interval and that there is little overlap between the set of cookies for which queries are received across successive days. In other words requests for the same cookie are highly localized in time making it hard to train a model based on historical data.").

[797] Launch Doc, "Online Reserve Price Optimization Launch Doc" (Sep. 14, 2017), GOOG-DOJ-13211589, at -589 ("Online prediction is accomplished by maintaining a sliding window of the N (current[ly] set to [redacted] most recent bids for each buyer-cookie pair in the high bid range [redacted]) and then applying a prediction function on the set of bids in this window in order to determine per-buyer reserve prices for the query at hand.").

[798] Launch Doc, "Online Reserve Price Optimization Launch Doc" (Sep. 14, 2017), GOOG-DOJ-13211589, at -590 ("The maximum of all reserves including those computed via other RPO models is then applied to the buyer's new bid in the auction.").

[799] Email from [redacted] to [redacted], "[Launch 124105] Per-buyer dynamic reserve price optimization on AdX - full launch [redacted] annual revenue from RTB buyers)" (Nov. 13, 2014), GOOG-DOJ-15419945, at -945 ("Experiments show [redacted] revenue for AdX buyers on AdX pubs and [redacted] for AdX buyers on AdSense pubs. Revenue impact on GDN is neutral."); "AdX Dynamic Price: Sell Side Review fact sheet" (Dec. 2014), GOOG-DOJ-13200158, at -158 ("Based on simulations and live experiments, AdX Dynamic Price increases revenue from AdX buyers (RTB + Static) by about [redacted]. GDN revenue and ROI is largely unaffected due to it being exempt. Benefit to publishers. Publishers benefit from this feature by earning increased revenue from AdX. In aggregate, revenue of AdX publishers increases by [redacted], and revenue for AdSense publishers increases by [redacted]%, for a network-wide increase of [redacted]%."); Presentation, "AdX Dynamic Price" (Dec. 11, 2014), GOOG-DOJ-13199910, at -930 ("AdX pubs [...] Total [redacted] (with GDN opt out").

revenues further: on the order of [redacted]% compared to no RPO.[800] The addition of Online RPO was estimated to further increase revenue by [redacted]% compared to the revenues obtained under the earlier RPO models.[801]

405. RPO was temporarily deactivated when AdX transitioned to the Unified First-Price Auction in September 2019.[802] Google redesigned RPO for application in the Unified First-Price Auction, including modifications to ensure that the same floor price was applied to each buyer within an auction (including Google Ads).[803] Optimized Pricing for the First-Price Auction was launched in June 2022, and it was enabled by default for all publishers using GAM (but it also included an opt-out for publishers).[804] In April 2022, Google also introduced an option (in beta release) for publishers to allow Google to

---

[800] Email from [redacted] to [redacted]m, "[Launch 132012] Cookie based Dynamic Reserve Price Optimization on AdX - full launch" (Sep. 29, 2015), GOOG-DOJ-15423462, at -462 [redacted]).

[801] Launch Doc, "Online Reserve Price Optimization Launch Doc" (Sep. 14, 2017), GOOG-DOJ-13211589, at -590 ("The launch candidate increases network-wide revenue by [redacted]%").

[802] "Reveal RPO floor to Adwords" (Nov. 11, 2019), GOOG-DOJ-14030931, at -931 ("Note that since moving to first-price auctions, AdX does not have any RPO floor yet (as of 2019/10/31) except small experiment traffic.").

[803] "Apply Dynamic Reserve Price to DFP Remnant Ads in First Price Auctions" (Jan. 5, 2020), GOOG-DOJ-14029750, at -750 ("And similar to how we are migrating legacy per-buyer publisher floor to unified pricing rules, we are making dynamic reserve price non-personalized, and applying that to both DFP remnant and backfill demand."); "Reveal RPO floor to Adwords" (Nov. 11, 2019), GOOG-DOJ-14030931, at -931 ("Previously in second price AdX auctions, Adwords was exempt from RPO floor [...]. This is no longer the case after AdX moved to first-price auctions. RPO floor applies to all demand in first-price AdX auctions.").

[804] "2022 Google Ad Manager releases archive," Google Ad Manager Help, https://support.google.com/admanager/answer/11586212?hl=en#zippy=%2Cjune-optimize-pricing-video-protections-troubleshooting-for-mcm-google-analytics-integration-webview-api-for-ads-updates-to-bid-rejection-reason ("June 6 Optimize pricing [...] Optimize pricing to reflect inventory's value[:] Optimized pricing increases auction floor prices to more accurately reflect and protect your inventory's value. Optimized pricing is enabled by default, but can be disabled via your network settings.").

automatically optimize the floor prices for slices of inventory in its Unified Pricing Rules (UPR) feature (which I discuss further in Section XII).[805]

D. Responding to Plaintiffs' Allegations

1. Google's Communications About RPO and the AdX Auction Format Were Not Misleading

406. Plaintiffs and their experts object to Google's communications about RPO to advertisers and publishers, with Professor Pathak characterizing RPO as a "secret auction manipulation program[]" that "reduced transparency for [Google's] customers."[806] But all of the Plaintiffs' experts overlook the fact that Google had flagged to its customers as early as 2014 the possibility of an optimization like RPO that modified a publisher's floor prices on some impressions.[807] Moreover, RPO was officially announced to the public on

[805] "Optimize floor prices in unified pricing rules (Beta)," Google Ad Manager Help, https://support.google.com/admanager/answer/11385824 ("Optimized floor prices (Beta) are available as a pricing option within unified pricing rules, in addition to fixed floor prices and target CPM. [...] You specify a slice of inventory in a unified pricing rule where Google will automatically optimize floor prices."); "2022 Google Ad Manager releases archive," Google Ad Manager Help, https://support.google.com/admanager/answer/11586212?sjid=14339833919693691538-NA#zippy=%2Capril-audience-segment-forecasting-optimize-floors-in-uprs-desktop-anchor-ads-linked-account-changes-mediation-chain-update-view-top-pricing-rules-atp-for-lgpd-update ("April 25 [...] Optimize floors in UPRs").

[806] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 174 ("Google's secret auction manipulation programs reduced transparency for its customers. Google leveraged its ability to reduce transparency when launching programs such as Bernanke, Dynamic Revenue Sharing (including v1, v2, and tDRS), and Reserve Price Optimization ('RPO')."). *See also* Fourth Amended Complaint ¶ 335 ("With RPO, Google abused advertisers' trust and secretly used their true value bids against them."); Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 283 ("Google concealed vital information from advertisers by concealing RPO."); Expert Report of J. Andrien (Jun. 7, 2024), at ¶ 35 ("Google also impacted advertiser behavior through its second-price auction representation and concealment of RPO.").

[807] "Ad Exchange auction model" (Aug. 24, 2014), GOOG-AT-MDL-C-000035250, at -250 ("The Google DoubleClick Ad Exchange may run limited experiments designed to optimize the auction. These experiments may include modifying the standard auction model or mechanics (e.g., a tiered, rather than second price auction); simulating ad calls and auctions; modifying the min CPM set by the publisher for an impression or otherwise adjusting publisher settings; or discounting certain bids submitted by buyers or otherwise modifying the priority of the bids submitted by buyers, in an effort to optimize the auction. Publisher's buyer/advertiser blocks will not be modified.").

May 12, 2016, roughly a year after the program was launched,[808] so any potential impacts from the alleged "concealment" of RPO would have been short-lived.

407. Plaintiffs allege that RPO misled advertisers because it "meant that the auction did not operate as a sealed second price auction as Google had advertised."[809] This is wrong. As I discuss in Section III.C.3, a second-price auction is the sealed-bid auction process that assigns the impression to the highest bidder for a price equal to the larger of the second-highest bid or the highest applicable floor price, set before bids are received. With or without a program like RPO in place, a second-price auction remains a second-price auction, as noted by Plaintiffs' experts Professors Weinberg and Pathak.[810] RPO never unsealed bids received in an auction to set the floor price for that auction.[811]

[808] Jonathan Bellack, "Smarter optimizations to support a healthier programmatic market," Google Ad Manager (May 12, 2016), https://blog.google/products/admanager/smarter-optimizations-to-suppor/ ("In our experiments to date, we have applied optimized pricing to about 15% of transactions, creating over 5% lift in revenue for publishers using the Open Auction. As we expand our experiments with optimized pricing, we will monitor its performance to ensure advertisers continue to get great ROI."). *See also* Email from [redacted] to [redacted]., "[ANNOUNCED] Smarter optimizations for DoubleClick Ad Exchange" (May 12, 2016), GOOG-DOJ-04934481, at -481 ("On Thursday afternoon we announced new optimizations for DoubleClick Ad Exchange, Optimized Private Auctions and optimized pricing in the Open Auction (aka RPO).").

[809] Fourth Amended Complaint ¶ 532. *See also* Expert Report of J. Andrien (Jun. 7, 2024), at ¶ 35 ("Google also impacted advertiser behavior through its second-price auction representation and concealment of RPO. Namely, I understand that Google's representation that it was running a second-price auction encouraged advertisers to bid their true value for impressions, which over time caused later AdX reserve prices to increase, which, in turn, led to a payoff loss for advertisers by decreasing win rates and increasing the average clearing price in later AdX auctions."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 348 ("With Reserve Price Optimization (RPO), advertisers believed they were in a standard second-price auction, but Google set artificially high reserve prices to optimize AdX revenue, often higher than the publisher's floor price.").

[810] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 285 ("Note that each individual auction in isolation is still a second-price auction with reserve [...] .") (describing an example of the operation of RPO); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 191 ("As a result, even though AdX ran a second-price auction at the time [...] .") (discussing RPO's effects on advertiser bidding behavior).

[811] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -159 ("[T]he price a buyer pays is not related to the bid in the present auction."), -161 ("[O]nly historical bids are analyzed - we do not use the bids in the auction to set the floor.").

408. Plaintiffs and some of their experts claim that Google's auction is not a "true" second-price auction by noting that a repeated auction process may fail to incentivize bidders to bid their true values.[812] Plaintiffs quote a Google employee who made similar errors: "Doesn't that undermine the whole idea of second price auctions?"[813] While any second-price auction is bidder-truthful—meaning that the bidder cannot increase its payoff *in that auction* by bidding anything other than its value for the impression—it has long been understood (and was noted by Professor Weinberg[814]) that a bidder might gain by untruthful bidding in a *repeated* second-price auction. For example, I wrote in a 2006 paper (referencing a 1990 paper by Rothkopf, Teisberg and Kahn[815]) that "[b]idders [in a second-price auction] may rationally be reluctant to report their true values, fearing that

[812] *See* Fourth Amended Complaint ¶ 346 ("All the while, Google continued to lead publishers and advertisers to believe that AdX operated a second-price auction, inducing advertisers to submit a sealed bid reflecting their true value."), ¶ 330 ("Consequently, when bidding into AdX and revealing their true value, advertisers relied on Google's representations that AdX was truly a sealed second-price auction."), ¶ 533 ("[...] because of RPO, Google's AdX auctions were not conducted as a true sealed second price auction as Google advertised. Rather with RPO, winning advertisers did not pay the true second price in an auction. Instead, they paid an artificially inflated price floor which Google set using advertiser's own bidding information against them. Second, Google encouraged advertisers to provide their true value in connection with what they thought was a sealed second price auction. The benefit of an advertiser bidding their true value into a second price auction is only realized when the auction is won, and the winner pays a second price below their true value bid. Instead, Google failed to disclose that it would use historical bidding information to artificially drive up the second price, by increasing publisher's preset price floor and replacing it with an artificially inflated floor derived from advertiser's historical bidding information."); Expert Report of J. Andrien (Jun. 7, 2024), at ¶ 35 ("Google also impacted advertiser behavior through its second-price auction representation and concealment of RPO. Namely, I understand that Google's representation that it was running a second-price auction encouraged advertisers to bid their true value for impressions, which over time caused later AdX reserve prices to increase, which, in turn, led to a payoff loss for advertisers by decreasing win rates and increasing the average clearing price in later AdX auctions."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 348 ("With Reserve Price Optimization (RPO), advertisers believed they were in a standard second-price auction, but Google set artificially high reserve prices to optimize AdX revenue, often higher than the publisher's floor price.").

[813] Fourth Amended Complaint ¶¶ 347, 537; Email from M. Loubser to D. Bradstock, "Re: [drx-prm] Re: SBS Australia meeting notes - April 30th" (May 1, 2015), GOOG-TEX-00537828, at -829.

[814] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 52 ("Truthfulness only holds when viewing this auction in isolation, it does not necessarily hold when considering a series of auctions [...].").

[815] Rothkopf, M.H., Teisberg, T.J., & Kahn, E.P. (1990). Why are Vickrey auctions rare? *Journal of Political Economy*, *98*(1), 94-109.

the information they reveal will later be used against them."[816] This does not change a second-price auction into something else, but highlights that the possibility of learning is an *unavoidable* complication when auctions are run repeatedly and sellers learn from experience how to set price floors. Even in that case, however, using a second-price auction format "minimizes the need to 'game' the system," as noted by a Google employee (as quoted by Plaintiffs).[817] In repeated second-price auctions, the possibility of learning is the *only* reason for a bidder to adjust its bids away from its value for an impression. That possibility of bid adjustment in the repeated second-price auction context does not change the basic fact that each individual auction is second-price.

409. Plaintiffs allege that Google was "misleading[]" when it claimed that RPO "helps advertisers by increasing the amount of inventory available for purchase programmatically."[818] But encouraging publishers to make more inventory available to programmatic sales is a likely result of *any* program that increases publisher yield per impression. This is a simple consequence of the well-known law of supply: all else equal, an increase in the price received by a seller increases the quantity supplied.[819]

410. Plaintiffs allege that—prior to the public announcement of RPO in May 2016—Google "continued to mislead publishers by encouraging them to adjust Google exchange floors in their publisher ad server [...] leading them to believe that they could control outcomes

[816] Ausubel, L.M., & Milgrom, P. (2006). The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham & R. Steinberg (Eds.), *Combinatorial Auctions* (pp. 17-40). MIT Press.

[817] Fourth Amended Complaint ¶ 531 (quoting AdExchanger, "Google's Scott Spencer On DoubleClick Ad Exchange Auction And Data Management," AdExchanger (Feb. 9, 2010), https://www.adexchanger.com/ad-exchange-news/googles-scott-spencer-on-doubleclick-ad-exchange-auction-and-data-management/.

[818] Fourth Amended Complaint ¶ 345.

[819] Mas-Colell, A., Whinston, M.D., & Green, J.R. (1995). *Microeconomic theory.* Oxford University Press, at 138.

and optimize yield through floors."[820] But publishers could still set effective floor prices when RPO was in place: the publisher would never receive a payment lower than the floor price it set. This means that publishers could always effect a floor price in AdX higher than the RPO floor. Even before RPO was introduced, EDA or the value CPMs of remnant line items could increase the effective floor prices for advertisers to be above the floor prices set by publishers. Moreover, RPO was a program with *very large* benefits to publishers: Per-Buyer RPO and Cookie-Based RPO together increased publisher revenues on the order of [redacted]%, and Online RPO increased publisher revenues even further.[821]

2. Publishers Benefited From RPO

411. Professor Weinberg claims that there are "reasons why some publishers ***might*** prefer outcomes without RPO than with RPO."[822] To support these claims, he provides examples

---

[820] Fourth Amended Complaint ¶ 342.

[821] Email from [redacted] to [redacted]m, "[Launch 124105] Per-buyer dynamic reserve price optimization on AdX - full launch [redacted] annual revenue from RTB buyers)" (Nov. 13, 2014), GOOG-DOJ-15419945, at -945 ("Experiments show [redacted]% revenue for AdX buyers on AdX pubs and [redacted]% for AdX buyers on AdSense pubs. Revenue impact on GDN [redacted]."); "AdX Dynamic Price: Sell Side Review fact sheet" (Dec. 2014), GOOG-DOJ-13200158, at -158 ("[E]xperiments, AdX Dynamic Price increases revenue from AdX buyers (RTB + Static) by about [redacted]%. GDN revenue and ROI is largely unaffected due to it being exempt. Benefit to publishers. Publishers benefit from this feature by earning increased revenue from AdX. In aggregate, revenue of AdX publishers increases by [redacted] and revenue for AdSense publishers increases by [redacted]%, for a network-wide increase of [redacted]%."); Presentation, "AdX Dynamic Price" (Dec. 11, 2014), GOOG-DOJ-13199910, at -930 ("AdX pubs [...] Total [redacted]% (with GDN opt out)"; Email from [redacted] to [redacted], "[Launch 132012] Cookie based Dynamic Reserve Price Optimization on AdX - full launch" (Sep. 29, 2015), GOOG-DOJ-15423462, at -462 ("We've previously experimented [with per-buyer RPO]. Experiments show [redacted]% revenue from AdX buyers across AdX and AdSense publishers incremental to [per-buyer RPO]."); Email from [redacted] to [redacted], "Re: RPO rollout schedule" (Sep. 21, 2015), GOOG-DOJ-15423317, at -319 ("The cookie pricing model alone generates about [redacted]% increase in adx buyer revenue compared to no dynamic pricing. When used together with the primary ('inventory') model, it generates [redacted]% revenue increase on top of the primary model ([redacted]% total increase)."); Presentation, "AdX Dynamic Price V2" (May 26, 2015), GOOG-DOJ-14000011, at-012 ("Inventory RPO impact: [redacted]% on AdX buyers [...] Cookie RPO impact: about the same (and less than expected from sim) [...] Combined impact: [redacted]% (almost additive) on adx buyers"); Launch Doc, "Online Reserve Price Optimization Launch Doc" (Sep. 14, 2017), GOOG-DOJ-13211589, at -590 ("The launch candidate increases network-wide revenue by [redacted]%").

[822] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 280 (emphasis added).

of publisher objectives that he claims are not supported by RPO, including publishers who are interested in "maximiz[ing] the number of sold ads (as opposed to maximizing their revenue)"[823] and publishers who "prefer[] to trust their own optimization over Google's."[824] But RPO was designed to support publishers "to earn the most money possible, with the least complexity,"[825] and the simplicity of its automation together with the significant increases in publisher revenues that Google measured in its RPO experiments (discussed in Paragraph 404 above) show that the program was very successful in achieving those goals. None of the Plaintiffs' experts provide actual evidence of publishers with the preferences that they speculate "might" exist or point to examples of publishers claiming to be harmed by Google's RPO program. Nor am I aware of any such evidence.

3. Publishers and Advertisers Were Not Harmed By "Concealment" of RPO

412. Plaintiffs and their experts allege that advertisers and publishers were harmed by the "concealment" of RPO.[826] Yet, as explained in Section II.B.3 of this report, processes for

[823] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 280.

[824] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 279.

[825] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -154 ("Our goal has always been to help publishers thrive and create sustainable businesses with advertising: to earn the most money possible, with the least complexity, all while providing users the best experience.").

[826] *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 272 ("The negative effects of RPO to advertiser payoff, and possibly some publishers' revenues, at least partially stem from Google's concealment of the conduct."); Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 191 ("As a result, by concealing RPO, Google interfered with the publishers' revenue maximization [...] As a result, both the publishers and the advertisers could act accordingly if they knew RPO was in effect. This potentially reduces their revenue or payoff compared to what they could have been."); Expert Report of J. Chandler (Jun. 7, 2024), at ¶ 360 ("It is my opinion that those undisclosed Google rule changes were contrary to the expectations of the auction participants and made it impossible for auction participants and competing exchanges to understand the rules that governed and applied to auctions run by Google, skewing decision-making and outcomes.") (discussing RPO, DRS v1, and Bernanke); Expert Report of J. Andrien at ¶¶ 34-35 ("For example, by concealing RPO, Google prevented publishers from effectively optimizing revenue, [...] Google also impacted advertiser behavior through its second-price auction representation and concealment of RPO.").

setting reserve prices, like those for computing bids, are routinely kept secret to avoid other participants exploiting the details of those processes to their own advantage, at the expense of the publisher. The concealment of a program like RPO serves the publisher's interests and should be expected by advertisers.

413. Professor Weinberg argues that a lack of transparency about the functioning of RPO "prevented the publishers from effectively optimizing their revenues,"[827] and Professor Pathak echoes these claims, alleging that "since the program was not announced to publishers, publishers would not be able to maximize their revenues under the scheme."[828] In reality, RPO was a service that helped publishers that had set some floor prices too low by raising those floor prices to increase their revenues, and Google experiments found that RPO had very large benefits for publishers. RPO did not prevent publishers from running experiments that could identify further improvements or even revenue-maximizing floor prices. Also, if a publisher had chosen the revenue-maximizing floor price for an impression, RPO would make *no change* to the floor price applied for that publisher. These same observations undermine Professor Weinberg's additional claim that "even after RPO [was] disclosed, publishers would still face challenges setting optimal reserves under RPO."[829]

414. Professor Weinberg alleges that "[a]dvertisers would change their bidding behavior had Google revealed RPO,"[830] and Professor Pathak agrees, alleging that "had they known

[827] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 279.

[828] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 190.

[829] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 279.

[830] Expert Report of M. Weinberg (Jun. 7, 2024), at Section IX.B.1.

[about RPO], the advertisers might have wanted to shade their bids."[831] But this incentive is not a result of programs like RPO or the many similar programs employed by non-Google supply-side intermediaries.[832] As noted above, in *any* auction setting involving repeated interactions, a buyer or seller needs to account for the possibility that another auction participant might learn from its past behavior and use that information in future interactions.[833] Even before the introduction of Optimized Pricing in April 2015, publishers were already adjusting price floors based on historical bids,[834] which suggests that RPO was automating an optimization function that publishers were already doing themselves. Plaintiffs' experts provide no evidence that the disclosure of the program in 2016 led to any change in bidder behavior (as would be expected under their theory), so

[831] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 191. *See also* Fourth Amended Complaint ¶¶ 532-33 ("Instead, advertisers were forced to pay significantly more for an ad they otherwise would have paid in the absence of RPO. [...] Had advertisers known that Google was manipulating publisher floors in this manner, they would have engaged in alternative bid strategies that did not disclose their true value for each impression.").

[832] [redacted] Rubicon Project, "Maintaining the Equilibrium: How Dynamic Price Floors Preserve the Integrity of the Automated Advertising Ecosystem" (2014), GOOG-AT-MDL-011234683, at -696 ("Rubicon Project gives sellers the opportunity to employ a sophisticated Dynamic Price Floor algorithm which works above a seller's hard floors."); [redacted]

[833] This is sometimes known as the "ratchet effect" in economic theory. *See* Freixas, X., Guesnerie, R., & Tirole, J. (1985). Planning under incomplete information and the ratchet effect. *The Review of Economic Studies*, *52*(2), 173-91; Bergemann, D., & Välimäki, J. (2019). Dynamic mechanism design: An introduction. *Journal of Economic Literature*, *57*(2), 235-74.

[834] Comms Doc, "Optimized pricing in the Open Auction Comms" (Mar. 23, 2018), GOOG-DOJ-04937154, at -155 ("Optimized pricing in the Open Auction automates the post-auction analysis and floor price updates that publishers are already doing and takes it a step further.").

the suggestion that the alleged concealment of RPO affected advertiser behavior is mere speculation.

XII. SELL-SIDE DYNAMIC REVENUE SHARING: INCREASING MATCH RATES AND PUBLISHER REVENUES

A. Overview

415. Dynamic Revenue Sharing (DRS) was a sell-side feature on AdX "that increases publisher and Google revenue by dynamically changing the AdX sell-side revenue share so that more auctions end with a winning buyer."[835] Although DRS evolved over time, each version involved reducing AdX's revenue share applied to some individual impressions to allow publishers to sell more impressions and AdX bidders to win more impressions. While DRS varied AdX's revenue share impression-by-impression, the average revenue share for each publisher remained close to and no lower than the contracted level throughout all three versions.[836] By maintaining or increasing the average revenue share received by the publisher, DRS could increase Google's profits only when it also increased the total revenues it paid to publishers.

416. Plaintiffs and their experts allege that "[a]lthough DRS could increase the volume transacted by publishers, it could also harm them due to the opportunity cost of clearing a bid at a lower price."[837] Although it is correct that one of the key benefits of DRS was an increased volume of transactions cleared, in their analysis of harms to publishers, Plaintiffs and their experts fail to account correctly for strategic adjustments: in the

[835] Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -321.

[836] The revenue share was closer to [redacted] during DRS v1. *See* Email from [redacted] "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -462 ("AdX margin: [redacted]"); Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -324 ("DRS v2 […] maintain[s] a Google share of [redacted]); Design Doc, "Truthful DRS Design Doc" (Mar. 24, 2017), GOOG-DOJ-13227256, at -261 ("keep the average AdX revshare at the contracted value [redacted].").

[837] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 839. *See also* Fourth Amended Complaint ¶ 325.

presence of DRS, publishers would be incentivized to increase floor prices. After accounting correctly for those adjustments, my analysis shows that DRS would likely increase total publisher revenues. These findings are consistent with Google's internal experiments at the time DRS was introduced.[838]

417. Plaintiffs' experts claim that buyers and publishers were unable to optimize their bids and floor prices because Google "misrepresent[ed] the sealed second-price auction."[839] This claim is incorrect for two reasons. *First*, starting from at least August 2015—before the launch of DRS v1—Google publicly disclosed on the AdX Help Center page that AdX could adjust its revenue share on individual impressions.[840] *Second*, buyers and publishers

[838] Presentation, "Overall Pub Yield with DRS(v2)" (Apr. 7, 2016), GOOG-DOJ-13235100, at -101 to -102 states a [redacted] in total publisher revenues when comparing pre-DRS to DRS v2 [redacted] in publisher revenue"; "that is the lift of DRS v2 (half-way with buy/pub side recollection) compared with no-DRS, since all the numbers in the deck are DRSv2 vs no-DRS"); Email from [redacted] to [redacted] al, "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -462 states an [redacted] [redacted] publisher revenues from AdX when comparing pre-DRS to DRS v1 ("AdX publisher payout: [redacted]").

[839] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 855. *See also* Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 188; Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 220.

[840] For the Help Center page before the launch of DRS v1, see Help Center Page, "Ad Exchange auction model" (Aug. 4, 2015), GOOG-AT-MDL-C-000035251, at -251 ("DoubleClick Ad Exchange determines the winning bidder based on the highest net bid submitted. Note that the net bid reflects any adjustments Ad Exchange may, at its discretion, have made to the bid submitted by the buyer for the purpose of optimizing the auction. [...] In some cases, the auction may close at a price lower than the reserve price applied, due to auction optimizations. Sellers are paid the Ad Exchange closing price, net of Google's revenue share, but will receive, subject to the terms governing their use of Ad Exchange, no less than the min CPM applied to the auction."). For the Help Center page before the launch of DRS v2, see Help Center Page, "Ad Exchange auction model" (Jun. 14, 2016), GOOG-AT-MDL-C-000035252, at -252 ("DoubleClick Ad Exchange determines the winning bidder based on the highest net bid submitted. Note that the net bid reflects any adjustments Ad Exchange may, at its discretion, have made to the bid submitted by the buyer for the purpose of optimizing the auction. [...] To optimize the auction, Google may choose to close an auction at a price lower than the reserve price that would have otherwise been applied. In such cases, the winning buyer may pay a price below the reserve and therefore receive a discount on its bid. A buyer that has received discount(s) on its bid(s) may face higher reserve prices in subsequent transactions to offset such discount(s). Subject to the terms governing their use of Ad Exchange, sellers are paid the Ad Exchange closing price, net of Google's revenue share, but will receive no less than the min CPM they specified for the auction. Unless the 'per-query revenue share' setting is enabled by a Seller, auction optimizations may result in an auction closing at a price lower than the reserve price that would have otherwise been applied. Because the Seller will always be paid at least its specified min CPM, the Seller may receive more than its contracted revenue share on the transaction. In subsequent transactions, the Seller's revenue share may then be reduced to offset the prior earnings in excess of the contracted revenue share, but the Seller will always receive at least its contracted revenue share across all its Ad Exchange transactions in a given month.").

can experiment to improve their own outcomes regardless of what information they have about the auction.

418. Plaintiffs argue that "DRS was exclusionary and inflicted significant harm on competition in the exchange market [...] Only Google's exchange could set its take rate on an impression-basis after peeking at all of its rival's net bids."[841] Plaintiffs' experts similarly claim that "DRS is most harmful to competition"[842] and that DRS "increases Google's win rate and revenue and further decreases the win rate and revenue of other exchanges."[843] In fact, DRS applied predominantly to allow the sale of impressions where the publisher-set floor prices were too high to be met, rather than ones for which floor prices were determined by header bidding.[844] "Peeking ahead at other exchanges' net bids"[845] was irrelevant in those cases, and DRS was output-expanding, allowing the sale of impressions that may not otherwise have been allocated[846] and increasing AdX's win rate and publisher revenues without affecting the win rates or revenues of those other exchanges. In other cases, which were fewer in number, DRS allowed the publisher to increase the price it received for an impression above its best offer from header bidding (via the so-called "last look"). DRS was an improvement in AdX's product offering that

[841] Fourth Amended Complaint ¶ 330.

[842] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 807.

[843] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 224.

[844] *See* "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273. Internal Google experiments show that, of the impressions on which DRS lowered the AdX revenue share, [redacted] occurred when the payment was determined by a third-party bid. By contrast, [redacted] occurred when the payment was determined by a publisher floor price. *See id. See also* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 32 ("In most cases when DRS applied, the reserve price was set by the publisher-set floor price.").

[845] Fourth Amended Complaint ¶ 328.

[846] Or would have been allocated to remnant line items that paid less than the publisher-set floor price.

increased payments to publishers, reduced the number of unsold impressions (expanding output), and increased the total value of impressions won by AdX advertisers.

B. How Dynamically Adjusting Revenue Shares By Impression Can Create Value By Selling More Impressions and Increasing Publisher Revenues

419. An intermediary, such as AdX, with more accurate predictions than the seller about a buyer's value can sell more impressions and increase seller revenues (along with its own profits) by dynamically adjusting the revenue share it applies to individual impressions. Because the minimum bid required to win an impression depends on both the floor price *and* the revenue share applied to the impression, AdX could change the probability that an impression sold by changing its revenue share on an impression. The goal of increasing publisher revenue and enabling the sale of additional impressions is consistent with the description of the objectives for DRS in Google's internal documents: "DRS is an optimization feature that increases publisher and Google revenue by dynamically changing the AdX sell-side revenue share so that more auctions end with a winning buyer."[847]

420. Here is one example that illustrates the possible benefits of a program like DRS. Suppose a publisher is selling impressions to a potential buyer through an intermediary that charges a 20% revenue share. The buyer's value (in CPM) is $1.00 for 10% of impressions, $1.25 for 80% of impressions, and $1.50 for the remaining 10% of impressions, but the publisher does not know for any given impression what the buyer's value will be.[848]

[847] Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -321.

[848] After applying the 20% revenue share, the net CPMs are $0.80, $1.00, and $1.20, respectively.

421. First, I calculate the publisher's maximum revenues in the absence of a program like DRS. In that case, the publisher can calculate, for each possible floor price, the probability that the impression will sell and use that information to determine the floor price that maximizes its revenue. In the example, if the publisher chooses a floor price of $1, after accounting for the intermediary's 20% revenue share, the buyer must bid at least $1.25 to win the impression, and in that case, the impression sells 90% of the time, leading to an expected revenue of 90¢ for the publisher. This is more than it can expect to earn with any other floor price.[849]

422. Now, suppose that the intermediary can perfectly predict the values of the advertiser and use that information to increase publisher revenues, while maintaining its 20% revenue share on average. To maximize the publisher's revenue, the intermediary can adjust its revenue share on a per-impression basis. When it predicts that the buyer's CPM is $1, it can set its revenue share at 0%, allowing the impression to sell to the buyer at the publisher's floor price of $1, and when it predicts the buyer has a CPM of $1.50, it can set its revenue share at 33%, selling the impression to the buyer at a price of $1.50 and passing the net revenue of $1 on to the publisher. The result is that the advertiser is able to purchase more impressions (the percentage of impressions sold rises from 90% to 100%), the publisher increases its revenues (from $0.90 to $1 on average per impression), and the intermediary, which still has an average revenue share of 20%, also increases its profits.

[849] If the publisher's floor price is 80¢ (leading to a pre revenue share floor price for buyers of $1), the impression sells 100% of the time, giving the publisher expected revenue of 80¢. If the publisher's floor price is $1.20 (leading to a pre revenue share floor price for buyers of $1.50), the impression sells only 10% of the time, giving the publisher expected revenues of 12¢. Any higher floor price for the publisher never sells the impression and any lower floor price leads to lower revenues.

423. This example is not contrived. It shows one way in which a better-informed intermediary can dynamically set its per-impression revenue shares to help a publisher sell more impressions, increasing its revenue. Similar effects can be achieved for very general distributions of advertiser values and different qualities of information available to the intermediary.[850]

424. This example is not intended to illustrate the exact operation of any one version of DRS, which used more complicated rules to determine revenue shares than that described in the example. The key properties that the example and DRS have in common is that when DRS observed or predicted that the buyer had a low valuation, it would lower the AdX revenue share on those impressions to allow additional sales, and that, in some versions, it would raise its revenue share on other impressions to maintain a fixed average AdX revenue share. The details of DRS evolved through three versions over time, and the way Google collected payment for its services also varied, but just as in the example, Google designed DRS to increase publisher revenues and the quantity of impressions sold, thereby raising its own revenue.[851]

C. DRS Increased Match Rates and Publisher Revenues and Evolved to Simplify Bidding for Advertisers

425. Google recognized the potential for a program like DRS to increase its profits while helping publishers sell more impressions and increasing their revenues, but it took time to

[850] This example is inspired by Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, *105*(3), 921-57. It deviates from their exact model, in which the better-informed party sets the (floor) price rather than the revenue share.

[851] *See* Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -321 ("DRS is an optimization feature that increases publisher and Google revenue by dynamically changing the AdX sell-side revenue share so that more auctions end with a winning buyer.").

engineer a program that did so effectively while maintaining its fixed average revenue share and simple bidding for advertisers. As shown in Figure 13, an excerpt from Google's launch documents for DRS v1, Google planned and implemented a gradual introduction of DRS, testing elements of its design, while monitoring the responses of advertisers and publishers and tracking the impacts of the program on publisher and advertiser outcomes.[852] Google launched DRS v1 expecting to "learn from a simple version and see responses."[853]

[852] *See* Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -958, -960 ("Holdback plan to assess long term buyer and seller response").

[853] Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -956.

Figure 13: Google's planned evolution of DRS[854]

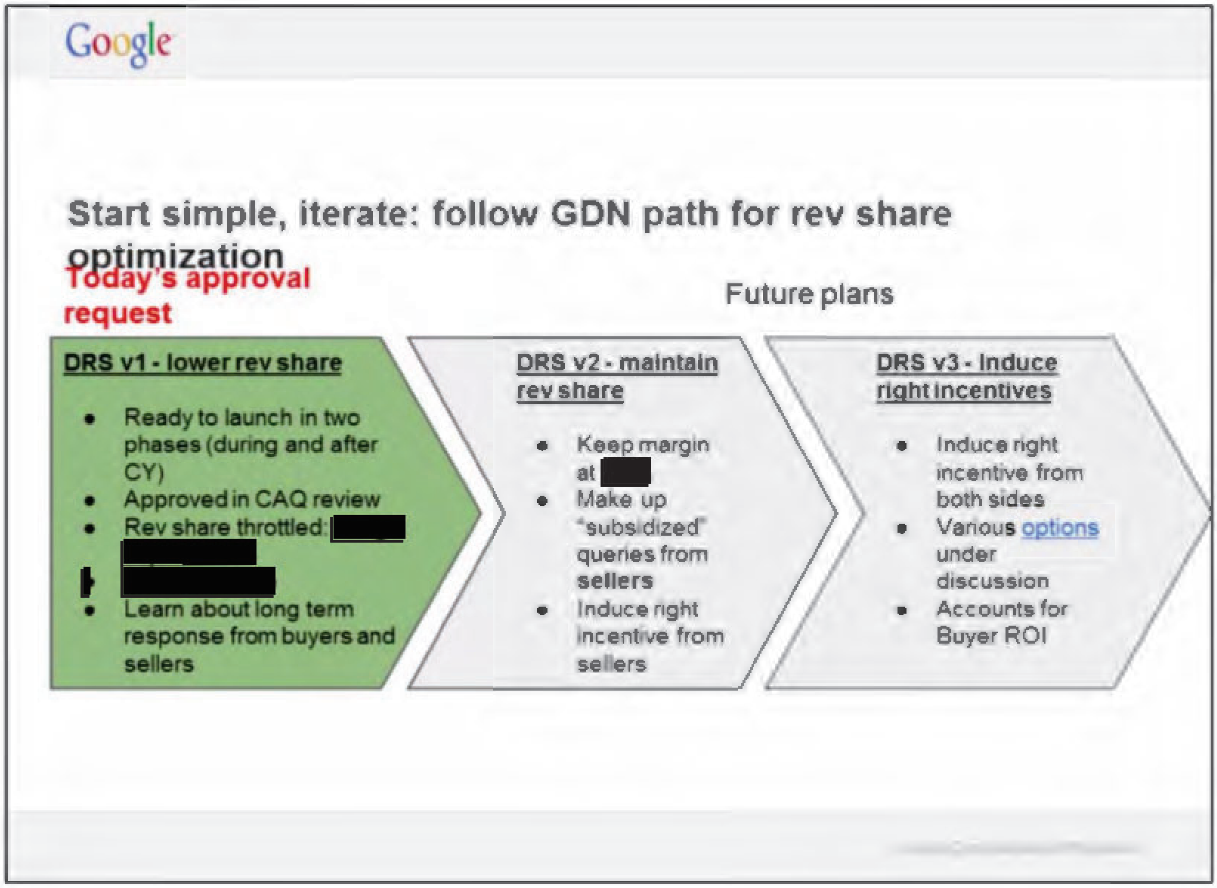

1. DRS v1 Lowered Per-Impression Revenue Shares to Sell More Impressions

426. In DRS v1, launched in August 2015, AdX made only one change to its auction rules: it reduced its revenue share on some impressions to allow more impressions to be sold.[855] Under DRS v1, AdX would first determine for each impression if the highest bid it had

[854] Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -958.

[855] Email from [redacted], "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -461 ("DRS clears queries when the highest bid is above the publisher floor, but not quite enough above the floor to cover the 20% AdX revenue share. In these cases we lower the revenue share per query as needed to increase transaction volume and increase match rate."); "AdX dynamic sell-side rev share (DRS v1) - project description / mini PRD" (Aug. 2014), GOOG-DOJ-03619484, at -484 ("Reduce the 20% rev share when there is no winner at 20% and an opportunity to find a winner with a reasonable, lower rev share").

received was in the "dynamic region," meaning that the bid would not clear the floor price if AdX charged its standard revenue share on that impression, but the bid would clear the floor price if AdX selected a 0% share.[856] On an impression for which the highest bid fell in this dynamic region, AdX would sometimes reduce its revenue share on that impression as needed to allow the bid to clear the floor price. AdX did not reduce its revenue share on each such impression: instead it "throttled" the application of DRS, meaning that it applied it probabilistically, adjusting the probability it applied DRS over time to ensure that the average AdX revenue share on impressions sold by each publisher and to each advertiser both did not drop significantly below the contracted AdX revenue share (for example, for publishers who had contracted an AdX revenue share of 20%, AdX throttled DRS v1 to maintain an average revenue share of at least 19%).[857]

427. Figure 14 displays an example of how DRS v1 could increase publisher revenues and the quantity of impressions transacted. In this example, the publisher sets a floor price of $1.00. Without DRS, AdX collects its contracted revenue share on each impression (in

[856] In its initial design of DRS, Google was agnostic about the source of the floor price that AdX faced. It could have been a floor price set by the publisher, a value CPM from a remnant line item (including header bidding line items), or a floor price GAM had calculated under its Reserve Price Optimization (RPO) program. DRS treated these floor prices the same way, without distinguishing their source. In July 2017, Google updated the DRS algorithm to exclude RPO floor prices, effectively ensuring that a fixed revenue share was applied to impressions for which the floor price was determined by RPO. *See* Launch Details Spreadsheet, Launch 193573 (Sep. 1, 2023), GOOG-AT-MDL-009644409, at cell B2 ("Disable dynamic revenue sharing for RPO prices").

[857] Email from M. Loubser to drx-pm@google.com et al., "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -461 ("DRS clears queries when the highest bid is above the publisher floor, but not quite enough above the floor to cover the 20% AdX revenue share. In these cases we lower the revenue share per query as needed to increase transaction volume and increase match rate. We limit how often we reduce the margin to maintain a >=19% average margin[.]"); Presentation, "Dynamic Sell-Side Revshare[:] GDN/DRX Summit 2015" (Nov. 2, 2015), GOOG-DOJ-13202659, at -670 ("Probabilistically throttle DRS[:] For an incoming query, we flip a coin to throttle queries (publisher) or winner (advertiser) when AdX margin is squeezed below a predefined threshold. […] Adjust throttling probabilities based on the difference in measured margin and pre-defined threshold."); Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -954 ("Buyers throttled based on whether query will clear using DRS, then checking buyer throttling table for winning buyer[;] Sellers throttled based on lookup of pub throttling state by query (query identifies pub)").

this example, 20%), which means that the impression can only sell if the advertiser bids at least $1.25.[858] This could leave revenue on the table for the publisher: if the advertiser bids $1.11, it does not win the impression, even though it is willing to pay more than the publisher's floor price for the impression. Under DRS v1, AdX would sometimes reduce its revenue share on the impression as needed to allow the impression to sell. If DRS v1 was applied to the same impression with a floor price of $1 and bid of $1.11, AdX would reduce its revenue share to 10%, allowing the bidder to win the impression while paying its bid of $1.11, and passing on $1.00 of that to the publisher, leaving 11¢ in revenue for AdX (equivalent to a 10% revenue share on that impression).[859, 860]

[858] To calculate the floor price facing the advertiser, divide the publisher's floor price by the publisher's revenue share (which is one minus the AdX revenue share). In this case, this calculation obtains $1.00/(1-0.2)=$1.25.

[859] Note that Figure 14—reproduced from an internal Google document—incorrectly calculates a 9% revenue share under DRS v1, but the actual clearing revenue share is closer to 10%.

[860] Note that 11¢ equals 10% of $1.11.

Figure 14: An example of DRS v1 adjusting revenue shares at an impression level to increase impressions transacted.[861]

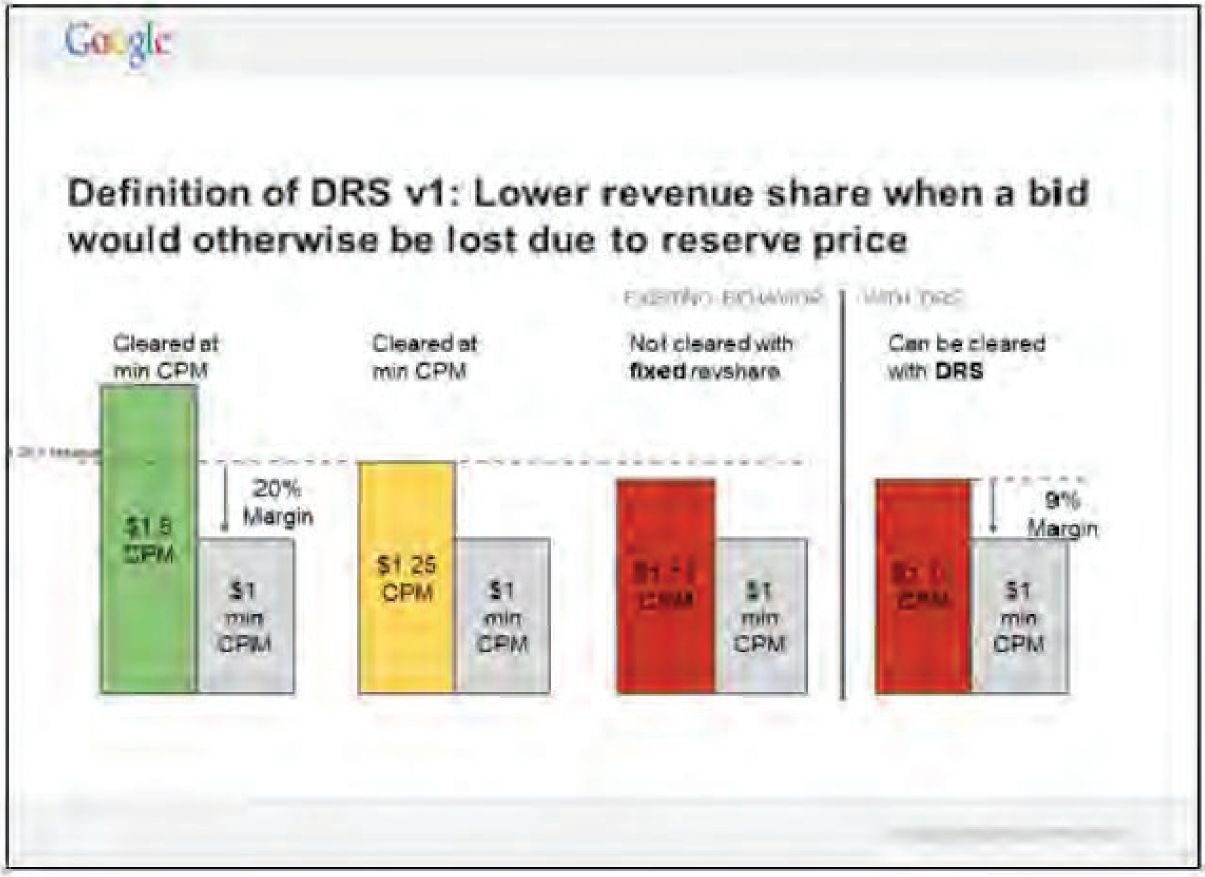

428. After the introduction of DRS v1, if publishers and bidders left their floor prices and bids unchanged, then both sides would benefit from the program: publishers would sell more impressions on AdX, increasing their revenues, and advertisers would win more impressions, increasing their surplus. I formalize this result in the following theorem, proved in Section XV.E.1.

[861] Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -954.

429. **Theorem 4**: If publishers did not change their floor prices and bidders did not change their bids, DRS v1 could only increase the number of impressions sold, publisher revenues, and advertiser surplus.

430. Theorem 4 assumes that publishers and bidders leave their floors and bids unchanged in response to DRS v1, which may be most relevant for a short-run analysis of the program. But, as I now explain, DRS v1 created new incentives for publishers and bidders to adapt their strategies to further increase their payoffs.

431. Under DRS v1, bidders were incentivized to reduce their bids for impressions. To see this, note that on each additional impression an advertiser won as a consequence of DRS v1, it was charged exactly its bid for the impression. This is not a threshold pricing rule (see Paragraph 61) because a winning bidder could sometimes lower the price it paid by reducing its bid for the impression. At the same time, publishers were incentivized to change their floor prices. Holding bids fixed, DRS v1 made higher floor prices less costly for a publisher, since AdX would sometimes compensate when the publisher set its floor price too high.

432. Incorporating all of these effects in a standard model of auction theory, I show in Theorem 5 that, if bidders and publishers adjust their bids and floor prices optimally, then publisher revenues increase under DRS v1 and advertiser surplus is approximately unchanged compared to the absence of DRS v1. The proof of Theorem 5 is in Section XV.E.2.

433. **Theorem 5:** Suppose that a publisher is selling an impression to a fixed set of bidders on AdX. The publisher does not know each bidder's value for the impression, and bidders do

not know each other's values for the impression, but all participants have the following information:

a. Each bidder knows its own value for the impression.

b. Each bidder's value is drawn from a commonly-known probability distribution and is statistically independent from other bidders' values.[862]

c. Each bidder determines a bid as a function of its value to maximize its surplus from the impression, given its probabilistic assessments about the bids of other bidders.

Then, if the publisher chooses revenue-maximizing floor prices, it earns a higher expected revenue on an impression to which DRS v1 is applied than it would without the program, and advertisers' surplus is unchanged.

434. Theorem 5 characterizes outcomes under the "independent private values model," which is a model adopted by Plaintiffs' experts.[863] While characterizing the theoretical effect of DRS v1 under more general modeling assumptions is difficult, there is a good reason to believe that it increased publisher revenues. As I show in Theorem 6, because DRS v1 reduces AdX's average revenue share, it can only be profitable to AdX if it increases publisher revenues from AdX. The proof of Theorem 6 is in Section XV.E.3.

[862] This implies that each bidder and the publisher can make a probabilistic assessment about other bidders' values and that estimates of other bidders' values would not be changed upon learning one bidder's value.

[863] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...] and it is a sensible assumption to make").

435. **Theorem 6**: Revenue for AdX under DRS v1 increases *only if* publisher revenues from impressions sold via AdX also increase. Additionally, every percentage increase in revenue for AdX results in a proportionally larger increase in revenue for publishers.

436. Internal Google experiments confirm the theoretical predictions that DRS v1 increased publisher revenues, finding that publisher revenues from impressions sold via AdX increased by [redacted] after the introduction of the program.[864] Google also found that DRS v1 increased the overall AdX match rate [redacted][865] In the experiment, AdX exactly hit the [redacted] revenue share target set in the design of DRS v1.[866]

2. DRS v2: Restoring a Fixed Revenue Share by Averaging

437. DRS v2, launched in December 2016, restored the average AdX revenue share to the levels that applied prior to DRS.[867] It accomplished that by tracking the payments by each buyer to each publisher using "debt accounts" to ensure that, while the AdX revenue share applying to an individual impression could be higher or lower, AdX received its standard revenue share on average over all impressions.[868]

---

[864] Email from [redacted]., "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -462 ("AdX publisher payout: [redacted]").

[865] Email from [redacted]., "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -461 ("Overall match rate for AdX publishers increases by [redacted]").

[866] Email from [redacted] "LAUNCHED! AdX Dynamic Revenue Share (DRS)" (Sep. 2, 2015), GOOG-AT-MDL-B-001391461, at -462 ("AdX margin: [redacted]").

[867] "AdX Dynamic Revshare v2 Launch Plan" (Aug. 2016), GOOG-DOJ-13208074, at -074 ("Launch DRS v2[:] Launched on 12/7").

[868] *See* Launch Doc, "AdX Dynamic Revshare v2: Launch Doc" (Jan. 13, 2016), GOOG-DOJ-13207875, at -879 ("In DRS v2, we expand per-query margin range to [redacted] with an objective to keep average adx margin at [redacted] over queries."). This average revenue share was a per-publisher and per-buyer average ("DRS v2 is implemented by dynamically expanding the AdX revshare based on the debt accumulated for each buyer and each seller.").

438. These debt accounts worked to maintain the average revenue share as follows.

439. On some impressions for which the highest bid was in the dynamic region, DRS v2 would decrease the AdX revenue share to ensure the impression would sell. Rather than charging each winning buyer its bid on those impressions with a discounted revenue share, DRS v2 charged each winning buyer an amount between its bid and the publisher's chosen floor price. AdX would then add a "debt" to the buyer's debt account, equal to the "discount" it had applied to the buyer's bid to allow that impression to transact, which is the amount that the buyer would need to raise its bid to win with the standard per-impression revenue share.[869] AdX would also add a "debt" to the publisher's debt account equal to its "discount" on the AdX revenue share, the amount that its floor price would need to be lowered in order to sell the impression with the standard per-impression revenue share.[870]

440. On some impressions for which the highest bid was above the amount required to win the impression in the absence of DRS, AdX would collect its standard revenue share plus an additional payment to recoup debts previously accrued by publishers and advertisers under DRS v2.[871] It did so while still charging the winning buyer a total amount less than its bid and paying the publisher more than its floor.[872] A winning buyer would be charged

[869] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 [redacted]

[870] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 [redacted]

[871] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 ("We attempt to collect debt in a later query for [redacted], which we call the non-dynamic region.").

[872] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -357 ("[W]e are careful to pick buyer_collection and publisher_collection to preserve constraints such as […] pay to the

the standard clearing price for the impression, plus a payment for any debt it had accrued, where this additional payment was chosen so that the total price of the impression was less than the bidder's bid.[873] If the buyer had set its own clearing price (by submitting a nonzero second bid), any additional payment it made as a consequence of its second bid was deducted from the debt balance as well.[874] [redacted].[875] I give a more detailed mathematical description of the workings of DRS v2 in Section XV.E.4.

441. Note that when recouping debts, AdX pays a portion of the recouped buyer debt back to the publisher. These amounts were chosen exactly to ensure that there is no double-charging of debt and, assuming all advertiser debts are recouped, the net debt accruing on average to publishers is zero, as described in Lemma 1, proved in Section XV.E.5.

442. **Lemma 1:** Suppose that AdX recoups all debts under DRS v2. Then, buyers accrue the full debt on each impression (equal to the difference between the floor price that would apply in the absence of DRS and the price it pays under DRS v2) and, after accounting

publisher at least his reserve[.] […] [T]he revised 2sided DRS never increases the price set above max(2nd bid, reserve).").

[873] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 [redacted]

[874] Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 ("Finally we consider a second way in which buyers can decrease their debt [redacted]

[875] Launch Doc, "AdX Dynamic Revshare v2: Launch Doc" (Jan. 13, 2016), GOOG-DOJ-13207875, at -879 [redacted]).

for the payment of buyer debt to publishers, publishers accrue zero debt on net in expectation.

443. As Lemma 1 shows, publishers on average accrue *no net debt* on an impression cleared by DRS v2, while buyers pay on average a net debt equal to the difference between the floor price that would apply in the absence of DRS and the price it pays under DRS v2. This means that, after accounting for later debts paid, the effective price of an impression cleared by DRS v2 for the buyer is equal to the floor price that would apply in the absence of DRS v2. Publishers can increase their total revenues as a result of DRS v2 without changing their floor prices, as shown in Theorem 7, proved in Section XV.E.6.

444. **Theorem 7:** If publishers do not change their floor prices and buyers do not change their bids, then DRS v2 can only increase the total number of impressions sold and total publisher revenues compared to the absence of DRS.[876]

445. As in DRS v1, the implementation of DRS v2 changed the incentives for publishers and buyers in the auction. In Theorem 8, I show that when buyers and publishers respectively adapt their bids and floor prices to DRS v2, auction theory predicts that buyer surplus and publisher revenues will be the same as in the absence of DRS. This happens because, given the way that debts were recouped under DRS v2, buyers who bid into the dynamic region end up paying the same on average as they would need to bid to win the impression in the absence of DRS. Theorem 8 is proved in Section XV.E.7.

[876] I assume that AdX performs enough transactions that all debts in DRS v2 are resolved.

446. **Theorem 8**: If buyers and publishers set bids and floors to maximize their payoffs after the introduction of DRS v2, then buyer surplus and publisher revenues are the same as in the absence of DRS.

447. These theoretical results assume that all debts accrued during DRS v2 are collected, or equivalently, that all agents choose bids and floor prices *expecting* that all debts are collected. In their initial internal experiments on DRS v2, Google engineers found that the vast majority of buyers and publishers had very small quantities of uncollected debts ([redacted], respectively).[877] If small amounts of debt were left uncollected under DRS v2, I would expect the results above to be quantitatively similar.

448. In experiments conducted after the launch of DRS v2, Google found that DRS v2 led to a [redacted] net increase in *total* publisher revenues compared to no DRS, including revenue from remnant demand.[878] This latter experiment was performed on [redacted] of impressions, and found that [redacted] publishers saw an increase in total revenue.[879] Revenue also increased for publishers that used header bidding.[880] In pre-launch experiments, Google found that publisher revenues were approximately unchanged between DRS v1 and DRS v2.[881]

[877] *See* Presentation, "Two-sided Dynamic Revenue Sharing" (Nov. 28, 2014), GOOG-DOJ-13199584, at -598.

[878] Presentation, "Overall Pub Yield with DRS(v2)" (Apr. 7, 2016), GOOG-DOJ-13235100, at -101 [redacted]"), -102 ("[T]hat is the [redacted] compared with no-DRS, since all the numbers in the deck are DRSv2 vs no-DRS").

[879] Presentation, "Overall Pub Yield with DRS(v2)" (Apr. 7, 2016), GOOG-DOJ-13235100, at -106 ("Currently running [redacted] experiment."), -108 to -110 ("[redacted] publishers make more money in aggregate.").

[880] *See* Presentation, "Overall Pub Yield with DRS(v2)" (Apr. 7, 2016), GOOG-DOJ-13235100, at -114.

[881] Launch Doc, "AdX Dynamic Revshare v2: Launch Doc" (Jan. 13, 2016), GOOG-DOJ-13207875, at -875 ("Compared to v1, we observe […] almost neutral publisher payout").

449. Starting with the launch of DRS v2, publishers that did not want to participate could opt out and buyers that did not want to participate could choose not to submit bids below the floor price that they received in the bid request from AdX.[882]

3. Truthful DRS: Simplifying Bidding While Maintaining the Benefits of Dynamic Revenue Sharing

450. The final iteration of DRS, called Truthful DRS or tDRS, launched in July 2018.[883] tDRS maintained the benefits of dynamically adjusting the AdX revenue share by impression, while implementing a threshold pricing rule for buyers, which simplified bidding.[884] Under tDRS, AdX would choose its base revenue share for each impression [redacted] *before* collecting bids from AdX bidders, so that a winning buyer's bid did not affect its price.

---

[882] Buyers that did not wish to participate in DRS could opt out by not bidding in the dynamic region. Since 2016, all bidders see the highest among all the floor prices, including those determined by the value CPMs of remnant line items. With this information, bidders could decide whether they wanted to submit bids in the dynamic region. *See* Launch Doc, "Including Third-Party Threshold in the Revealed Reserve Prices to AdX Buyers" (Aug. 9, 2016), GOOG-DOJ-13208800, at -800 ("[redacted]."). Publishers could opt out of DRS in the user interface. *See* Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -326 ("You may choose to opt-out of revenue share based optimizations in the AdX UI. If you opt-out we will apply your contracted revenue share to every Open Auction query and you will not benefit from the increased revenue from this optimization.").

[883] "2018 Sellside Launches Revenue Evaluation" (Jul. 19, 2019), GOOG-DOJ-13949282, at tab "Q3Q4 2018," row 7.

[884] Another source of complexity under the first two versions of DRS was that, to optimize bids, a buyer would need to track performance across many auctions (ones on which revenue shares were discounted and others for which debt were repaid), making experiments by that buyer on subsets of impressions more difficult. This was noted by Google internally as another motivation for the transition to tDRS. *See* "The Future of DRS + RPO" (Jul. 2, 2016), GOOG-DOJ-13205869, at -869 ("[T]his mechanism [DRS] can make it challenging for buyers who adopt sophisticated bidding strategies which rely on learning the smallest bid with which they could have won the auction.").

451. Like DRS v2, tDRS maintained its average revenue share for each publisher. Truthful DRS accomplished this by maintaining debt accounts for each publisher.[885] Unlike DRS v2, tDRS did not target an average AdX revenue share by buyer. For each impression, AdX chose a revenue share based on its predictions of AdX buyers' bids.[886] [redacted].[888] In short, AdX would discount its per-impression revenue share when it *predicted* that doing so would allow the impression to sell, and AdX would reclaim the discounts on other impressions for which it *knew* that a higher revenue share would not prevent the impression from selling. I provide the mathematical details of the operation of tDRS in Section XV.E.8.

452. There are several important properties of the tDRS program:

[885] This is different from the approach under DRS v2, in which debts were effectively paid by advertisers (see Lemma 1 in Section XV.E.5).

[886] Design Doc, "Truthful DRS Design Doc" (Mar. 24, 2017), GOOG-DOJ-13227256, at -260 [redacted]

[887] Design Doc, "Truthful DRS Design Doc" (Mar. 24, 2017), GOOG-DOJ-13227256, at -261 ("[redacted]

[888] Design Doc, "Truthful DRS Design Doc" (Mar. 24, 2017), GOOG-DOJ-13227256, at -261 [redacted].").

a. *For any fixed floor price, more impressions would be sold with tDRS than without it:* To see this, note that, without tDRS, AdX always applied its contracted revenue share to each impression, whereas with tDRS, AdX discounted its revenue share on some impressions, making it more likely for those impressions to sell. On the other hand, AdX increased its revenue share on impressions only if it already knew there was a sufficiently high bid for that impression, so that this increase in revenue share never reduced the number of impressions sold.

b. *tDRS was bidder-truthful:* tDRS used a threshold pricing rule, meaning that the price paid by a winning bidder was the lowest bid that they could have made to win the impression. In any single auction, that property makes it optimal for a buyer to bid its value, simplifying the bidding problem for buyers.

c. *On average, each publisher would receive its contracted share of the revenue:* The average revenue share on each impression for publishers, after accounting for adjustments to the debt account, was equal to the contracted revenue share: wherever the revenue share on an impression was below the standard rate, the publisher would repay debt on future impressions to restore the contracted average revenue share; wherever the revenue share was higher than the standard rate, the publisher had received an offsetting discount on another impression.

453. If publishers do not change their floor prices in response to tDRS, these three properties imply that tDRS would *always* increase publisher revenues from AdX compared to no DRS. This follows since bidding incentives were unchanged and more impressions were sold, so that the total revenue collected from bidders must increase, which, since the same

average revenue share applied to the publisher, means that the publisher's total revenue must also increase.

454. Similarly to DRS v1, tDRS reduced the cost to publishers associated with setting a higher floor price, creating an incentive for a publisher to increase its floor prices in response to tDRS. If a publisher also used header bidding, tDRS created an incentive for the publisher to further "inflate" the header bid it reported to GAM (that is, to trigger a line item in GAM with a value CPM greater than the header bidding bid, see Section X.D.2). Truthful DRS strengthened the incentive to inflate header bids because, via that strategy, a publisher could ensure that an impression sold on AdX only if the *net* payment it would receive (after accounting for debts it incurred under tDRS) was larger than the publisher's expected payments from header bidding. Google's internal documents suggest that many publishers did inflate header bids, as predicted.[889]

455. After accounting for these incentives to increase floor prices and inflate header bids, I show in Theorem 9 that a publisher can guarantee that tDRS increases its *total* revenues

[889] *See* Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -506 ("We've anecdotally heard from some publishers that they inflate the value CPM of remnant line items [...] publishers used to do this even before HB was popular."); Presentation, "First-price bidding" (Aug. 12, 2019), GOOG-DOJ-11406673, at -677 ("How boost works[:] The publisher inflates the HB bid before sending it as a floor to AdX[.] This is done to increase Adwords cost and to provide a better comparison between Adwords and header bidder bids[.]"); Email from [redacted], "Unified Auction Changes (Sellside) Executive Update - Aug 12, 2019" (Aug. 13, 2019), GOOG-DOJ-09713317, at -319 ("Today, these [publisher-]inflated CPMs are used to provide price pressure for AdX […] In practice, […] many publishers […] apply a boost to Header Bidding bids"); Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -777 ("Last look [...] incentivizes pubs to inflate ('boost') the floor sent to AdX"). *See also* [redacted]

from all sources, including revenue from impressions with a competing header bidding offer. I provide a proof of Theorem 9 in Section XV.E.9.

456. **Theorem 9**: If publishers adjust their floor prices on AdX to maximize profits after the introduction of tDRS and tDRS accurately predicts buyers' bids, total publisher revenues from all demand sources will be higher with tDRS than with a fixed revenue share.

D. Responding to Plaintiffs' and Their Experts' Allegations

457. The opinions provided by Plaintiffs' experts are almost entirely limited to initial versions of DRS that were always intended to be improved upon. As I discussed above, Google planned a gradual introduction of DRS, testing elements of its design while monitoring the effects on advertisers and publishers, and some changes in publisher and advertiser outcomes could be expected as Google improved its DRS designs.[890]

1. Plaintiffs' and Their Experts Provide Misleading Analyses of DRS's Effects

458. Plaintiffs' experts allege that DRS harmed non-Google exchanges. Professor Gans states "[i]n combination with DA, DRS is most harmful to competition. [...] In a Waterfall setup, [...] DRS may foreclose rival exchanges from an opportunity to bid on those impressions."[891] Professor Weinberg concludes that DRS "increases Google's win rate and revenue and further decreases the win rate and revenue of other exchanges."[892] The evidence that I have reviewed suggests that these arguments omit DRS's main effects, which align with the program's stated purpose.

[890] *See, e.g.*, Presentation, "AdX DRS v1 launch review" (Feb. 13, 2015), GOOG-DOJ-13199952, at -958.

[891] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 807.

[892] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 224.

459. Recall that DRS was intended to "[increase] publisher and Google revenue by dynamically changing the AdX sell-side revenue share so that more auctions end with a winning buyer" rather than going unsold.[893] Internal Google experiments show that the most common effect of DRS was to cause an otherwise unmatched impression to be sold, rather than to win additional inventory off other exchanges.[894] Clearing these impressions would increase AdX's win rate and publisher revenues, but would do so *without* affecting the win rates or revenues of those other exchanges.

460. [redacted], DRS still provided publishers with the opportunity to increase their revenues beyond what non-Google exchanges offered by allowing impressions to sell with "inflated" header bids in GAM.

461. Plaintiffs' experts allege that DRS harmed publishers by decreasing their revenue. Professor Gans focuses on DRS v2, stating "[t]he combined effects of Last Look and DRS v2 led to revenue losses for the publishers for the following reasons. When DRS v2 led to a decrease in the take rate to clear a binding price floor generated by Last Look, which is equal to the Header Bidding winning bid, the impression is awarded to AdX; the

893 Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -321.

894 *See* "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273. Experimental results show that, of the impressions on which DRS lowered the AdX revenue share, [redacted] occurred when the payment was determined by a third-party bid. By contrast, [redacted] occurred when the payment was determined by a publisher floor price. *See id.*

895 *See* "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273. Experimental results show that, of the impressions on which DRS lowered the AdX revenue share, [redacted] occurred when the payment was determined by a third-party bid. By contrast, [redacted] occurred when the payment was determined by a publisher floor price. *See id.*

publisher gets paid only 1 cent higher than the Header Bidding winning bid."[896] To support his conclusion, he provides an example in which the highest header bid is $6 and the two highest bids from AdX are $7 and $5, and AdX's reserve is $3. Without DRS, AdX would apply a [redacted] revenue share, resulting in bids of [redacted] and [redacted] net of the revenue share. In this case, [redacted] wins the impression and the publisher receives [redacted]. Under DRS v2, AdX would apply a [redacted] revenue share, resulting in a high bid of [redacted]. In this case, AdX wins the impression and the publisher receives [redacted]. However, he claims the publisher would be worse off because "the publisher would incur the debt of [redacted]."[897, 898]

462. Professor Gans' analysis is flawed on three counts: (1) his calculation of debt under DRS v2 is misleading; (2) he omits the more frequent case in which the floor price is one set by the publisher; and (3) he omits the benefits when an optimizing publisher handles header bids correctly. While my discussion of (1) is specific to DRS v2, my discussion of (2) and (3) applies to all three versions of DRS.

463. *First*, by only looking at a single transaction, Professor Gans ignores the details of debt recollection under DRS v2. In his example, the advertiser also incurs a debt of $[redacted]. Assume that after his example, a second impression arrives. The two highest bids from AdX are $7 and $5 and AdX's reserve is $3, but there is no header bid. Under

---

[896] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 809.

[897] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 810.

[898] Professor Gans assumes that DRS v2 charged the winning AdX bidder its full bid of [redacted] hence the [redacted] revenue share. The version of DRS v2 that launched charged the winning AdX bidder the average of its bid and the floor price, equaling $[redacted] and leading to a revenue share of [redacted]%. The publisher would then incur debt equal to [redacted] However, this difference does not qualitatively affect the analysis, so I use his numbers for ease of exposition.

DRS v2, AdX wins the impression, and the bidder pays [redacted] The publisher receives [redacted] Both the bidder and the publisher's debt accounts are cleared after this transaction. In sum, the publisher receives the same revenue from the two impressions under DRS v2 as it would without DRS. More generally, I show in Lemma 1 that publishers accrue zero net debt in expectation, meaning that even though the publisher incurs [redacted] of "debt" in Professor Gans' example, the publisher receives the same revenue from AdX in expectation as it would receive from the header bidder. Plaintiffs' expert Professor Weinberg also reaches this same conclusion.[901]

464. *Second*, as explained above, Professor Gans' example applies only when the floor is set by a header bidding line item, but the more frequent effect of DRS was to create additional revenue for the publishers by helping it sell impressions that would otherwise go unsold.[902] Professor Gans does not attempt to weigh the revenue and efficiency gains from the additional transactions DRS enabled against its alleged downsides.

465. *Third*, Professor Gans' example would never arise if the publisher adopted a better strategy than directly passing the header bid to AdX. Such a publisher would *inflate* the header bid before reporting it to DFP, allowing the publisher to earn strictly more revenue from the sale of the impression. In Professor Gans' example, if the header bid is inflated

---

[899] The bidder's payment equals the clearing price plus the bidder's debt.

[900] The publisher's payment equals [redacted] the bidder's payment, minus the publisher's debt.

[901] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at Appendix G, ¶ 65.

[902] *See* "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273. Experimental results show that, of the impressions on which DRS lowered the AdX revenue share, [redacted] occurred when the payment was determined by a third-party bid. By contrast, [redacted] occurred when the payment was determined by a publisher floor price. *See id.*

to [redacted], then under DRS, AdX would win the impression only if it is able to provide a buyer willing to pay *more* than [redacted], thereby increasing the expected publisher revenue by at least [redacted] over the header bidder's [redacted] bid.

466. Google's internal experiments also showed that publishers benefited from DRS.[903] Citing a 2016 experiment in which DRS v2 led to revenue increases [redacted] publishers, Professor Gans focuses on the finding that "[redacted] saw their revenue decrease by around [redacted]."[904] That focus ignores the significant revenue increases for the other [redacted] publishers, which more than offset the minor revenue decreases experienced by the remaining [redacted] publishers. In addition, Professor Gans' interpretation overlooks that the document itself posits that "the fact that some publishers have their overall revenue decreased by DRS comes from noise due to the fact that this is a small experiment."[905] Finally, starting with the launch of DRS v2, publishers that did not want to participate in DRS could opt out and buyers could also avoid DRS by not submitting bids lower than the floor price shared in the bid request.[906]

[903] Presentation, "[redacted].

[904] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 840.

[905] "Does DRS make more money for publishers?" (Mar. 17, 2016), GOOG-DOJ-13204346, at -349.

[906] Buyers that did not wish to participate in DRS could opt out by not bidding in the dynamic region. Since 2016, all bidders see the highest among all the floor prices, including those determined by the value CPMs of remnant line items. With this information, bidders could decide whether they wanted to submit bids in the dynamic region. *See* Launch Doc, "Including Third-Party Threshold in the Revealed Reserve Prices to AdX Buyers" (Aug. 9, 2016), GOOG-DOJ-13208800, at -800 ("[redacted] Publishers could opt out of DRS in the user interface. *See* Comms Doc, "Dynamic Revenue Share" (Jan. 29, 2020), GOOG-DOJ-15130321, at -326 ("You may choose to opt-out of revenue share based

2. Google Did Not Mislead Buyers and Publishers About DRS

467. Plaintiffs' experts claim that buyers and publishers were not able to respond optimally to DRS because Google "concealed material information."[907] Professor Weinberg states, "Since publishers believed that AdX runs a regular second-price [auction] with their given reserve price and a [redacted] a strategic publisher would set a price floor that maximized their revenue under these circumstances. Had they known AdX dynamically adjusted its take rate, publishers would set different price floors."[908] Similarly, Professor Gans states that "Google intentionally omitted telling publishers that they were enrolled in DRS,"[909] and Professor Pathak states that "Google never disclosed this conduct to advertisers or publishers who sell or buy impressions through AdX."[910]

468. These claims of deception are contradicted by disclosures in the AdX Help Center starting from at least August 2015—before the launch of DRS v1. It says, "DoubleClick Ad Exchange determines the winning bidder based on the highest net bid submitted. Note that the net bid reflects any adjustments Ad Exchange may, at its discretion, have made to the bid submitted by the buyer for the purpose of optimizing the auction. [...] In some cases, the auction may close at a price lower than the reserve price applied, due to auction optimizations. Sellers are paid the Ad Exchange closing price, net of Google's revenue share, but will receive, subject to the terms governing their use of Ad Exchange, no less

optimizations in the AdX UI. If you opt-out we will apply your contracted revenue share to every Open Auction query and you will not benefit from the increased revenue from this optimization.").

[907] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 219.

[908] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 219.

[909] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 828.

[910] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 188.

than the min CPM applied to the auction."[911] Before the launch of DRS v2, the Help Center page was updated to include the possibility that some impressions would close above or below the contracted revenue share, while still ensuring that publishers would receive at least their contracted revenue share in a given month.[912,913]

469. Professor Weinberg's claim that "[s]ince DRSv1 was not disclosed to the advertisers, they would still bid their true value for impressions" is doubly wrong, both because DRS v1 was disclosed and because buyers could quickly and easily detect its application.[914] Every time DRS v1 affected a buyer's outcome, it caused the buyer to win an impression on which it had bid below the floor price and to pay the amount of its bid.[915] Without DRS, that outcome was impossible. DRS leaves an unmistakable trace that can be detected

---

[911] *See* Help Center Page, "Ad Exchange auction model" (Aug. 4, 2015), GOOG-AT-MDL-C-000035251, at -251 ("DoubleClick Ad Exchange determines the winning bidder based on the highest net bid submitted. Note that the net bid reflects any adjustments Ad Exchange may, at its discretion, have made to the bid submitted by the buyer for the purpose of optimizing the auction. [...] In some cases, the auction may close at a price lower than the reserve price applied, due to auction optimizations. Sellers are paid the Ad Exchange closing price, net of Google's revenue share, but will receive, subject to the terms governing their use of Ad Exchange, no less than the min CPM applied to the auction.").

[912] *See* Help Center Page, "Ad Exchange auction model" (Jun. 14, 2016), GOOG-AT-MDL-C-000035252, at -252 ("DoubleClick Ad Exchange determines the winning bidder based on the highest net bid submitted. Note that the net bid reflects any adjustments Ad Exchange may, at its discretion, have made to the bid submitted by the buyer for the purpose of optimizing the auction. [...] To optimize the auction, Google may choose to close an auction at a price lower than the reserve price that would have otherwise been applied. In such cases, the winning buyer may pay a price below the reserve and therefore receive a discount on its bid. A buyer that has received discount(s) on its bid(s) may face higher reserve prices in subsequent transactions to offset such discount(s). Subject to the terms governing their use of Ad Exchange, sellers are paid the Ad Exchange closing price, net of Google's revenue share, but will receive no less than the min CPM they specified for the auction. Unless the 'per-query revenue share' setting is enabled by a Seller, auction optimizations may result in an auction closing at a price lower than the reserve price that would have otherwise been applied. Because the Seller will always be paid at least its specified min CPM, the Seller may receive more than its contracted revenue share on the transaction. In subsequent transactions, the Seller's revenue share may then be reduced to offset the prior earnings in excess of the contracted revenue share, but the Seller will always receive at least its contracted revenue share across all its Ad Exchange transactions in a given month.").

[913] DRS v2 was launched in December 2016. "AdX Dynamic Revshare v2 Launch Plan" (Aug. 2016), GOOG-DOJ-13208074, at -074.

[914] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 220.

[915] DRS v1 lowers AdX's revenue share on impressions in which its highest bid net of the [redacted] revenue share is below the floor price.

every time it reduces AdX's revenue share. A minimally attentive buyer who failed to read Google's disclosure could still discover DRS by noticing when a bid below the floor price was accepted at a price equal to its bid and, by applying logic or just conducting routine experiments, optimize its bidding strategy to account for that possibility.

3. Non-Google Exchanges Also Dynamically Adjusted Revenue Shares

470. Professor Gans claims that only AdX had the requisite scale and incentive to implement DRS. He states, "Had Google not had monopoly power on the sell side of the market or been vertically integrated into the exchange market, it would not have had the ability or the incentive to undertake this conduct that harmed competition."[916] In fact, the same incentive applies to every exchange and the ability of other exchanges to undertake this conduct is revealed by the evidence. Non-Google exchanges [redacted] implemented similar programs, increasing and decreasing their revenue shares on individual impressions based on the floor prices and the bids they had received.[917]

[916] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 777.

[917] *See* [redacted] *See also* Sarah Sluis, "Explainer: More On The Widespread Fee Practice Behind The Guardian's Lawsuit Vs. Rubicon Project," AdExchanger (Mar. 30, 2017), https://www.adexchanger.com/ad-exchange-news/explainer-widespread-fee-practice-behind-guardians-lawsuit-vs-rubicon-project/ ("Exchanges might also vary their fee based on the difference between an advertiser's bid and the clearing price, PubMatic confirmed.").

4. DRS Did Not Remove Barriers to Low-Quality Ads

471. Professor Gans alleges that "DRS harmed publishers by lowering floors for low-quality ads."[918] Professor Gans provides no evidence of systematic differences in ad quality or that DRS actually resulted in low-quality ads winning auctions, instead asserting without support that "[p]ublishers use price floors to control the quality of ads showing on their webpage."[919] In reality, if a publisher wishes to exclude some types of ads, it can use content filters to achieve that result.[920] If the publisher wishes to show the ad only when the price is sufficiently high, it can set its floor price to accomplish that too: DRS allows the ad to transact only if the publisher receives at least that floor price.

5. Plaintiffs Mischaracterize Google's Combined Usage of DRS and RPO

472. Plaintiffs claim Google used DRS and RPO "in conjunction with each other to further the advantages provided by any one of these programs in isolation," stating "RPO would increase the price paid by an advertiser by raising the floor, while DRS would ensure that the advertiser's bid nevertheless cleared the higher floor set by RPO, which exacerbates harm to competition in the exchange market."[921] This depiction of Google's usage of DRS and RPO is misleading. Early on, Google saw "disab[ling] DRS on RPO floors" as a "[s]tep along the way" to "[m]aking DRS truthful,"[922] and eventually did disable DRS

918 Expert Report of J. Gans (Jun. 7, 2024), at ¶ 855.

919 Expert Report of J. Gans (Jun. 7, 2024), at ft. 1080.

920 *See* Google, "Block sensitive categories," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/2541069?sjid=14973969736947852707-EU#available-sensitive ("You can block groups of ads that are considered 'sensitive' due to the nature of the business or ad […] Our system classifies ads automatically, and we don't rely on advertiser-provided categorization.").

921 Fourth Amended Complaint ¶ 350.

922 Presentation, "AdX Dynamic Price" (Nov. 30, 2016), GOOG-DOJ-13206921, at -928.

on RPO-set floors in preparation for the launch of tDRS.[923] In practice, DRS and RPO interacted on less than [redacted] of impressions.[924]

[923] Launch Details Spreadsheet, Launch 193573 (Sep. 1, 2023), GOOG-AT-MDL-009644409, at cells B2, D2 [redacted]). tDRS was launched in July 2018. *See* "2018 Sellside Launches Revenue Evaluation" (Jul. 19, 2019), GOOG-DOJ-13949282, at tab "Q3Q4 2018," row 7.

[924] *See* "DRS and RPO interaction in Simulation" (Sep. 20, 2016), GOOG-AT-MDL-007375273, at -273. In the first table, the entry under the column "True" and the row "DYNAMIC_RESERVE" represents the percentage of impressions in which 1) AdX's bid fell in the dynamic region and 2) RPO set AdX's floor price, with a value of [redacted]

XIII. OPEN BIDDING AND UNIFIED FIRST PRICE AUCTION: IMPROVING PUBLISHERS' ABILITY TO ACCEPT BIDS FROM NON-GOOGLE EXCHANGES

A. Overview

473. Google began testing **Open Bidding** (previously known as Exchange Bidding, and internally also known as Jedi, EBDA, and demand syndication) in 2016[925] and officially launched the program in April 2018.[926] Open Bidding allowed publishers to accept bids for impressions from non-Google exchanges on Google Ad Manager, increasing the thickness of Google's platform. Accepting bids from other exchanges on GAM—in a way that protected the interests of both advertisers and publishers—required careful design and ultimately led to other changes on GAM, including the transition to the Unified First Price Auction (UFPA) in 2019.[927]

474. Google described Open Bidding as its "answer to header bidding."[928] Relative to header bidding, Open Bidding brought many advantages to publishers, including speed, streamlined payments, simpler configuration, and reduced computational burden on end

[925] *See* Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -348.

[926] *See* Jonathan Bellack, "Exchange Bidding now available to all customers using DoubleClick for Publishers," Google Ad Manager (Apr. 4, 2018), https://blog.google/products/admanager/exchange-bidding-now-available-to-a/ ("Exchange Bidding now available to all customers using DoubleClick for Publishers […] With Exchange Bidding, publishers can increase revenue by allowing multiple exchanges to compete with each other—and with DoubleClick Ad Exchange—in a unified auction.").

[927] Comms Doc, "Ad Manager Unified 1st Price Auction" (Sep. 27, 2019), GOOG-DOJ-09714662, at -662 ("[W]e are transitioning publisher inventory to a unified, 1st price auction for Google Ad Manager.").

[928] Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337 ("Demand Syndication is our answer to header bidding - a superior product for allowing pubs to get per-query bids from non-AdX exchanges[.]").

users' browsers.[929, 930, 931] The popularity of Open Bidding among publishers speaks to these benefits.[932]

475. Plaintiffs' allegations regarding Open Bidding and the UFPA are factually incorrect and fail to account for the incentives of participants.

a. *First*, Plaintiffs' allegations that Open Bidding was devised to "maintain its exchange monopoly and exclude competition"[933] from other exchanges using header bidding is inconsistent with the details of the Open Bidding program. Contrary to Plaintiffs' descriptions, publishers and exchanges were not forced to adopt Open Bidding and could maintain existing header bidding configurations. The data I have reviewed suggests that the vast majority of publishers using Open Bidding continue to use header bidding.

---

[929] In this section, I use the term "header bidding" to refer to client-side header bidding. Some use the term "header bidding" to include server-side header bidding tools, including Open Bidding. For a comparison of Open Bidding and header bidding, see Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -440 to -442.

[930] For a comparison of timeouts, see Comms Doc, "RTB Timeouts" (Oct. 2019), GOOG-DOJ-15232606, at -609 ("Google's lower bid timeouts should have a slightly better user experience with lower latency").

[931] *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24; Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -438 ("Eliminate operational inefficiencies such as line item complexity and latency that exists with header bidding […] Easy to set up, view/analyze reports and unified payments […] Allows exchanges to respond to RTB call-outs […] Provides integrated reporting and billing for exchange bidding transactions won by 3rd party exchanges"), -441 to -442.

[932] As of March 2023, [redacted] North American publishers using Google Ad Manager used Open Bidding for some impressions. Of those publishers, [redacted] use header bidding. This statistic is based on [redacted]" *See* [redacted], Appendix A (Oct. 6, 2023), GOOG-AT-MDL-C-000012826, at -874. Publishers are grouped by their network_id and ranked according to the number of queries matched in the dataset. Code for this result can be found in code/hb_monitor_ob_freq.py in my supporting materials, and the output is saved in code/logs/hb_monitor_ob_freq.txt.

[933] Fourth Amended Complaint ¶ 367.

b. *Second*, Plaintiffs and their experts allege that Google shared minimum bid to win (MBTW) data with Open Bidding exchanges (but not header bidding), allowing them "to adjust their future bidding strategy to continue trading ahead of exchanges returning bids through header bidding and underpaying for publishers' impressions"[934] and to "effectively recreate[] Last Look."[935] But bidders and bids vary from auction to auction, so Plaintiffs and their experts' suggestion that AdX buyers could exactly predict competitors' bids is incorrect. These allegations also significantly overstate the importance of receiving MBTW data and fail to consider numerous alternatives that enabled those other bidders to optimize bids.

c. *Third*, Plaintiffs and their experts claim that, in 2018, Google began redacting data fields in auction data shared with publishers to "cripple[] publishers' ability to measure the success of rival exchanges in header bidding."[936] But, these claims ignore the historical context of changes in data-sharing policies, which were

---

[934] Fourth Amended Complaint ¶ 380 ("Specifically, in 2019, DFP began sharing sensitive pricing information derived from publishers' sensitive clearing auction records (which Google called "Minimum Bid to Win" data) with exchanges in Exchange Bidding. Google's AdX exchange and other exchanges in Exchange Bidding use this data to adjust their future bidding strategy to continue trading ahead of exchanges returning bids through header bidding and underpaying for publishers' impressions.").

[935] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 199 ("Minimum Bid to Win data effectively recreates Last Look, because AdX and Exchange Bidding buyers can use the information from Minimum Bid to Win on their next auctions, while Header Bidding buyers cannot."). *See also* Fourth Amended Complaint ¶ 381 ("Google compounded this Exchange Bidding advantage with a new secret bid optimization scheme that allowed Google to recapture the advantages it had under Last Look. The new scheme uses information about publishers' ad server user IDs and rival exchanges' bids to accurately predict the amount to bid, effectively permitting Google to re-engineer the ability of AdX and Google's ad buying tools to trade ahead of rivals exchanges in Exchange Bidding. As a Google planning document outlines: 'If we knew our competitor's bid exactly, we can simply bid a cent above that[.] But we don't have this information before the auction, so we need to predict [the] competitor's bid.'").

[936] Fourth Amended Complaint, Section VII.D.3.iii ("Google cripples publishers' ability to measure the success of rival exchanges in header bidding (2018 to present)."). *See also* Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 150 ("The redaction of the DT protected AdX against the threat of Header Bidding because it removed publishers' ability to measure Header Bidding results and effectively target users."); Expert Report of J. Gans (Jun. 7, 2024), at ¶¶ 683-85 ("In short, by changing the way that KeyPart was generated (or "re- encoding" it), Google removed the ability for publishers to match DT files together. [...] By redacting these fields, Google prevented publishers from knowing relevant identifying information about auction winners.").

introduced in response to demands from buyer-customers and emerging privacy concerns. While Professor Gans asserts that "Google's claim that it redacted data based on privacy concerns is pretextual,"[937] he primarily justifies his opinion by a mischaracterized quote, ignores other evidence, and does not acknowledge Google's contractual obligations. Plaintiffs and their experts also fail to account for publishers' alternative means for measuring the "performance of exchanges in header bidding,"[938] such as experimentation and the use of other data fields.

B. Open Bidding: Google's Response to the Flaws of Header Bidding

476. Open Bidding was designed as Google's "answer to header bidding," incorporating more real-time bids from non-Google exchanges into Google's auction.[939] Although implementations of header bidding also enabled publishers to collect real-time bids from non-Google exchanges, those implementations tended to have several downsides, as described next.

---

[937] *See* Expert Report of J. Gans (Jun. 7, 2024), at Section VII.D.2 ("Google's claim that it redacted data based on privacy concerns is pretextual.").

[938] Fourth Amended Complaint ¶ 387.

[939] Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337 ("Demand Syndication is our answer to header bidding - a superior product for allowing pubs to get per-query bids from non-AdX exchanges.").

477. *First, latency.* Common implementations of header bidding increased latency.[940] An internal Google study found that header bidding increased ad loading times by approximately [redacted], and academic researchers found that header bidding latency was around [redacted] times that of the waterfall.[941] Latency not only degrades the end user experience, but in so doing, also harms advertisers and publishers.[942] By increasing the likelihood that an end user leaves the website before the ad is presented, latency harms publishers by potentially preventing a user from viewing an ad for which the publisher would be compensated if the ad had been displayed.[943] Latency also harms advertisers by reducing the effectiveness of online display advertising campaigns.

[940] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24 ("Most header bidding has traditionally taken place client-side, meaning the page sends out requests to individual ad exchanges and other demand sources, processes the responses, and then runs an auction, all via Javascript code running on the page. This may introduce latency issues and slow page loads."). *See also* Presentation, "Header Bidding T1 Impact" (Sep. 24, 2015), GOOG-DOJ-04430492, at -501 ("3. User's browser triggers the buyer's pixel, which communicates with the Header Bidding buyer's server. a. This may result in page load latency as the page will not load until the header script completes […] 4. HB buyer's server returns a decision […] (This results in further latency as the DFP tag is dependent on the Header Bidder's call resolving.)"),-502 ("Cons: Increased latency - especially in mobile and video = decreased user experience."); Vishveshwar Jatain, "Understanding Header Bidding And How To Leverage It," Forbes (Sep. 17, 2019), https://www.forbes.com/sites/forbescommunicationscouncil/2019/09/17/understanding-header-bidding-and-how-to-leverage-it/?sh=332097315c1 ("Client-side header bidding […] increased page latency because executing auctions takes bandwidth and computing resources.").

[941] Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -956 ("All things equal, there is a [redacted] in ads load speed because of Header Bidding"); Pachilakis, M., Papadopoulos, P., Markatos, E. P., & Kourtellis, N. (2019). No More Chasing Waterfalls: A Measurement Study of the Header Bidding Ad-ecosystem. In *IMC '19: Proceedings of the Internet Measurement Conference* (pp. 280-293).

[942] Interactive Advertising Bureau, "Glossary: Digital Media Buying & Planning" (Apr. 2016), https://www.iab.com/wp-content/uploads/2016/04/Glossary-Formatted.pdf, at p. 10 ("Latency sometimes leads to the user leaving the site prior to the opportunity to see the ad.").

[943] *See, e.g.*, DoubleVerify, "Latency in Digital Advertising: A Guide for Publishers," DV Publisher Insights (Oct. 21, 2019), https://pub.doubleverify.com/blog/latency-in-digital-advertising-a-guide-for-publishers/ ("High latency kills user experience and publisher revenue."); Google, "[UA] Latency and why it impacts Google Ads Clicks and Analytics Sessions," Google Analytics Help (accessed July. 24, 2024), https://support.google.com/analytics/answer/4589209?hl=en ("As a general rule, users on the internet are not very patient. This is evident from studies such as the KissMetrics study that elicited some sobering statements like: 'A one-second delay in page response can result in a 7% reduction in conversions.' along with '47% of consumers expect a web page to load in two seconds or less.' What does this mean for you? If your website loads too slowly, then there is a possibility that users are leaving and going to your competitors, especially if competitors are able to deliver the same content quickly.").

478. *Second, payment discrepancies.* There were reports of payment discrepancies, where publishers' expected receipts from header bidders did not match eventual payments.[944] In contrast to impressions sold on AdX (where AdX would act as a clearinghouse, collecting payments from bidders on behalf of publishers), publishers needed to manage their own billing and reconciliation for impressions sold via header bidding. Discrepancies of up to [redacted] could arise because different partners might apply different standards for verifying that impressions were not fraudulent or double-counted.[945]

479. *Third, complexity.* Header bidding could be challenging for a publisher to configure, requiring it to place code in the header of its website to request, collect, and compare bids before calling the publisher's ad server.[946] Header bidding also required a publisher to

---

[944] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 24 ("Header bidding is also not transparent because, although the publisher 'accepts' the impression at a certain price, the header bidder may not actually pay the sum indicated in its bid."); Presentation, "Header Bidding Observatory #1" (Jan. 2017), GOOG-DOJ-AT-01027937, at -955 ("A comparison of HB reports vs DFP reporting showed significant discrepancies [in revenue]"). *See also* James Curran, "For Publishers, Header Bidding Discrepancies Can Outweigh Revenue Lift," AdExchanger Opinion Page (Jul. 8, 2016), https://www.adexchanger.com/the-sell-sider/publishers-header-bidding-discrepancies-can-outweigh-revenue-lift/ ("Publishers need to create a more realistic calculation of header bidding revenue by factoring discrepancies into their line-item valuations. Some header bidding solutions can cause up to a [redacted] discrepancy between the publisher ad server impression reports and the impression reports from the programmatic partner. That means a [redacted] once you account for the adjustments made by the exchange for viewability, verification and performance tracking.").

[945] James Curran, "For Publishers, Header Bidding Discrepancies Can Outweigh Revenue Lift," AdExchanger Opinion Page (Jul. 8, 2016), https://www.adexchanger.com/the-sell-sider/publishers-header-bidding-discrepancies-can-outweigh-revenue-lift/ ("Publishers need to create a more realistic calculation of header bidding revenue by factoring discrepancies into their line-item valuations. Some header bidding solutions can cause up to a [redacted] between the publisher ad server impression reports and the impression reports from the programmatic partner. That means a [redacted] once you account for the adjustments made by the exchange for viewability, verification and performance tracking.").

[946] One publisher described the time required to integrate an SSP with a leading header bidding wrapper as "20 hours of work" (versus Open Bidding's "20 minutes"). Sarah Sluis, "Google Ad Manager Builds A Bridge To Prebid—But Don't Call It A Two-Way Street," AdExchanger (Apr. 27, 2022), https://www.adexchanger.com/platforms/google-ad-manager-builds-a-bridge-to-prebid-but-dont-call-it-a-two-way-street/. Another industry source compared header bidding to previous ad configurations: "It's not just a little more work, it's probably 100X as much work to traffic for most publishers." Ad Ops Insider, "Header Bidding Explained Step-by-Step" (Jun. 8, 2015), https://www.adopsinsider.com/header-bidding/header-bidding-step-by-step/.

take on many of the tasks that were traditionally performed by ad servers (or pay another publisher tool to do so), including billing, reporting, data security, and ad fraud prevention.[947]

480. Google found that publishers sought the ability to accept real-time bids directly from other exchanges on Google's platform.[948] In 2016, as header bidding was growing in popularity with publishers, Google started developing and testing Open Bidding, a product to accept real-time bids in that way—without the drawbacks of header bidding.[949]

C. Meeting the Challenge of an "Auction of Auctions": The Evolution of Open Bidding and the Unified First Price Auction

481. Designing an "auction of auctions" to improve on the flaws of header bidding while protecting the interests of both advertisers and publishers was challenging, in part because of differences between the auction formats on AdX and other exchanges. AdX ran a second-price auction, while many other exchanges had transitioned to

[947] *See* James Curran, "For Publishers, Header Bidding Discrepancies Can Outweigh Revenue Lift," AdExchanger Opinion Page (Jul. 8, 2016), https://www.adexchanger.com/the-sell-sider/publishers-header-bidding-discrepancies-can-outweigh-revenue-lift/ ("Managing header bidding discrepancies, like much of digital ad tech, is complex and time consuming. In order to make these changes, publishers need to recalibrate hundreds if not thousands of line items, which makes header bidding just as manual as a waterfall setup. This manual labor is another costly factor that publishers must include when determining the value of header bidding."); Ross Benes, "Unraveling header bidding's problems with user data," Digiday (Mar. 20, 2017), https://digiday.com/media/header-bidding-security/ ("[A]n overlooked downside [of header bidding] is that it can expose user data by allowing all bidders to access audience data. Header bidding also makes it easier for fraudsters to hide in the noise created by the vast amount of data points that come from multiple parties bidding on all available impressions. […] Hlavacek pointed out that with multiple partners bidding on all impressions available in the auction, header bidding significantly increases the amount of data points in the exchanges. Other sources said it's this data deluge that is most problematic when it comes to security.").

[948] Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337.

[949] *See, e.g.*, Presentation, "Demand Syndication" (Feb. 17, 2016), GOOG-DOJ-09459336, at -337.

non-second-price auction formats.[950] Bids that are optimized for different auction formats are not easily compared: in a first-price auction, bidders optimize by shading their bids below their values to maximize their payoffs, while bidders in a second-price auction find it optimal to bid their values, as I discuss in Section III.C.3. Designing an auction that incorporated bids from auctions run by other exchanges ultimately led to additional changes in the AdX auction design, including the transition to the Unified First Price Auction.

1. *"Alpha" Design Highlights Difficulty of Auction of Auctions*

482. The first "alpha" design of Open Bidding, in June 2016, incorporated bids from authorized exchanges in the AdX second-price auction, which AdX would compare to bids from Google's buy-side tools and Authorized Buyers.[951] When an auction was won by an Open Bidding buyer, GAM collected a revenue share of [redacted][952]

483. By continuing to receive bids from DSPs and other Authorized Buyers in a second-price auction, the Open Bidding design aligned with the way that those buyers had previously competed in the AdX second-price auction. Even so, this prototype design of Open Bidding was reportedly unpopular with non-Google exchanges because AdX could win

[950] Between late 2017 and early 2018, during the design phase of Open Bidding, Google estimated that the percentage of all queries conducted in "First-price auction" and "Second-price auction with anomalies" formats increased from [redacted] *See* Presentation, "DV360, Third Party Exchanges, and Outcome-Based Buying" (Oct. 16, 2018), GOOG-DOJ-12038253, at -298 ("In fact, exchanges are openly moving more and more to first price auctions.").

[951] *See* Presentation, "Exchange Bidding 'Last Look' leak" (Apr. 5, 2017), GOOG-DOJ-14024199, at -200 to -201 ("The initial design for exchange bidding in dynamic allocation (EBDA) worked almost exactly like Header Bidding when used with DFP. In the initial EBDA design, the bid submitted by an exchange was treated like a price derived from a line item and could become the floor price in the AdX auction. This part of the process was sometimes referred to as 'last look.'").

[952] Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 27 ("Up to at least December 2021, for publishers that utilized Open Bidding, when an auction was won by an Open Bidder, Google Ad Manager's standard charges for web display ads were [redacted] for GAM360 customers and [redacted] for other GAM customers, and Google's standard charges for app and instream video ads were [redacted].").

the auction with any bid higher than the bid made by each of the exchanges (and any other floor prices), at a price determined by the second-price auction logic (which could equal the largest bid from the other exchanges).[953] Although different in nature from the so-called "last look" over remnant line items (including header bids),[954] this Open Bidding design was also characterized as a "last look" for AdX.

484. By their nature, "alpha" launches are used to identify problems with a product's design before testing the product on a larger group of users. The challenges associated with the "alpha" launch of Open Bidding highlighted the complication associated with incorporating bids from other exchanges on AdX. Doing that effectively was not as simple as treating the bids from other exchanges in the same ways as other bids in the AdX auction. Instead, incorporating bids from other exchanges required additional changes to the AdX auction design.

2. Google Removes the So-Called "Last Look" Over Open Bidding, Challenging Its Second-Price Auction Design

485. In March 2017, Google responded to feedback and redesigned the Open Bidding auction to remove the so-called "last look" of AdX over other exchanges that participated in Open Bidding.[955] To accomplish that, Google *first* conducted a second-price auction among its AdX bidders with the same floor price it used before Open Bidding, and *then*

[953] *See* Sarah Sluis, "Google Removes Its 'Last-Look' Auction Advantage," AdExchanger (Mar. 31, 2017), https://www.adexchanger.com/platforms/google-removes-last-look-auction-advantage/ ("During a Google forum late last year, the biggest concern voiced by publishers and exchanges centered on this last-look advantage.").

[954] The main difference is that the value CPMs of remnant line items were not really bids for the impression—so that the publisher could choose reports to influence the AdX "last look," as discussed in Section X—whereas the bids made by exchanges under Open Bidding were actual bids for the impression that would later be used as the basis of payments.

[955] "2018 Sellside Launches Revenue Evaluation" (Jul. 19, 2019), GOOG-DOJ-13226855, at tab "Q1Q2 '17 Launch News," cell R60.

compared the AdX auction's clearing price to the best Open Bidding exchange bid on a first-price basis.[956, 957] A consequence of this procedure is that, for an AdX bidder to win the impression, AdX would need to have *two* bids (or a bid and a floor price) higher than the best first-price bid from other exchanges in Open Bidding.[958] This design was used for Open Bidding's "open beta" version, through the general release in April 2018, until the transition to the Unified First Price Auction in 2019.[959]

486. By maintaining the second-price auction on AdX, this design of Open Bidding was bidder-truthful (as defined in Section III.C.3.a) for bidders on AdX. But because the second-price auction's clearing price was not optimized as a bid into the first-price Open Bidding auction, AdX bidders were disadvantaged. To illustrate this disadvantage, suppose that AdX received two bids for an impression, $5 and $1, and the best first-price bid from another exchange was $1.20 (coming from a bidder with value of, say, $2). The auction outcome is determined by comparing AdX's clearing price from its internal second-price auction (that is, the second-highest bid of $1) to the best Open Bidding exchange bid (that is, $1.20), resulting in the $1.20 bid winning the auction. In this case,

956 That is, the AdX clearing price (net of AdX revenue share) would be compared to the best bid from Open Bidding (net of Open Bidding revenue share). If the net AdX clearing price was larger, the AdX bidder would win the impression and the publisher would receive the AdX clearing price; otherwise, the Open Bidding bidder would win and pay its bid.

957 *See* Presentation, "Exchange Bidding 'Last Look' leak" (Apr. 5, 2017), GOOG-DOJ-14024199, at -201 ("AdX and all EBDA SSPs run auctions and submit their bids[.] All bids are compared in a first price auction[.] SSPs are not acting as AdX 2nd price[.] Example[:] AdX highest bid- $5[.] AdX second bid- $2[.] SSP - 4$[.] Result: SSP wins at $4").

958 If there were no bids from AdX bidders above the floor price and that floor price was set by a non-guaranteed line item or the EDA price, AdX would submit the floor price as a bid into the first-price auction between exchanges, and would allocate the impression to the relevant non-guaranteed line item or direct deal if that bid won the Open Bidding auction.

959 *See* Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 26. Open Bidding's open beta launch was in May 2017. *See* Launch Details Spreadsheet, Launch 181852 (Aug. 31, 2023), GOOG-AT-MDL-009644401, at cell C2.

however, the AdX bidder willing to pay $5 is disadvantaged: it loses the auction despite having the highest value and bidding more than any other bidder. The publisher may also miss out on possible revenue improvements, because the highest bidder's willingness to pay much more has no effect on the final price in the Open Bidding auction.

487. A Google document summarized this problem as follows: "The clearing price from the internal AdX 2P auction[] may not always be competitive in the unified auction."[960] Internally, Google described this complicated auction design as "chaos."[961]

488. Despite these flaws, an internal Google presentation showed that the launch of Open Bidding resulted in approximately [redacted] increased revenue for publishers and was approximately [redacted] for Google.[962] Impressions sold by Open Bidding did not displace impressions sold by header bidding: instead, a Google analysis found that the share of impressions allocated to header bidding increased somewhat after the launch of Open Bidding.[963] Google's data showed that the main effect of Open Bidding was to replace impressions sold by *static* remnant bids with impressions sold by *real-time* bids in Open Bidding.[964] Having more impressions allocated using real-time bids improved platform thickness and efficiency. The introduction of Open Bidding was also celebrated

---

[960] Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -505.

[961] Presentation, "Welcome to Exchange Partnerships Executive Conference 2019" (May 2019), GOOG-DOJ-13501237, at -243.

[962] Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -761 ("[Open Bidding] was expected to [redacted]"), -771.

[963] Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -775 ("Median HB Win Rate per publisher [redacted]

[964] *See* Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -775 ("Publishers are moving from a waterfall based approach to RTB (EB and Header Bidding)").

by some competing exchanges, with OpenX reporting that "[e]xisting OpenX publisher partners who enabled [Open Bidding] through the OpenX Exchange experienced an average 48% increase in programmatic revenue from OpenX."[965]

489. Following the release of Open Bidding (which was still known externally as Exchange Bidding), Google launched Network Bidding in 2019.[966] Network Bidding offered "mega [ad] networks"[967] an additional alternative to traditional waterfall mediation, with AppLovin being the first network to participate.[968] Like Open Bidding, Network Bidding charged participants a [redacted] revenue share.[969] Bids from Network Bidding buyers competed in real-time against Open Bidding exchanges and AdX on a first-price basis.[970] If an ad network previously competed for inventory via static prices or the waterfall, I would expect real-time bids and the transition to Network Bidding to increase publisher revenues by the same logic that applies to Open Bidding above.

[965] OpenX, "Google & OpenX Release Study Showing Publisher Partners Experience 48% Revenue Lift Through Google Exchange Bidding Collaboration" (Feb. 15, 2018), https://www.openx.com/press-releases/google-openx-revenue-lift/.

[966] *See* Email from [redacted], "Re: AppLovin Open Bidding Live with First Beta Publisher!" (Jan. 10, 2019), GOOG-DOJ-AT-00849635, at -636 ("[W]e have AppLovin fully launched as our first Open Bidding network!").

[967] *See* Presentation, "FAN DFP & AdMob (Jedi 4 Networks) Opportunity" (Apr. 2017), GOOG-DOJ-09875881, at -889 ("Jedi for Networks * (NEW) Intended for [...] [redacted]").

[968] *See* Email from [redacted]l., "Re: AppLovin Open Bidding Live with First Beta Publisher!" (Jan. 10, 2019), GOOG-DOJ-AT-00849635, at -636 ("[W]e have AppLovin fully launched as our first Open Bidding network!").

[969] *See* Presentation, "Jedi For Networks" (Jan. 31, 2018), GOOG-DOJ-15226550, at -565 (table showing revenue shares for categories of demand and inventory, showing that Network Bidding had a [redacted] revenue share).

[970] *See* Product Requirements Doc, "PRD: Unified 1P auction and Pricing rules" (Jul. 25, 2018), GOOG-DOJ-03998505, at -505 ("AdX buyers (GDN, DBM and 3P RTB) submit second-price bids to AdX, meaning they compete in a 2P auction in which the winner pays the second-highest eligible bid or the price floor. With the introduction of Jedi (Exchange Bidding and Network Bidding), the winner of the second-price AdX auction submits the clearing price of the second price auction to another auction (unified auction), where it competes with EB and NB bids on a first price basis.").

3. Google's Transition to the Unified First Price Auction Further Increased Efficiency

490. To understand why Google transitioned to a first-price auction, consider the requirements it faced. Suppose that the ad allocation process must (1) employ an "auction of auctions" in which the winning exchange is the one with the highest auction clearing price; (2) ensure that the clearing price in its auction depends only on the bids in its auction and not on the bids in any other auction (avoiding what others have called a "last look"); and (3) ensure that the highest net bid by any advertiser wins. Applying simple logic to these three requirements implies that the clearing prices in each exchange can depend only on the highest bid in that exchange's auction, and not on the second-highest bid, so the requirements cannot be satisfied if any exchange uses a second-price auction. To achieve these requirements, *all participating exchanges—including AdX—must use first-price auctions*. For AdX, the switch from its historic second-price auction to a first-price auction was a major change that would present new challenges for all its bidders, because they would need to adapt their bidding procedures to the new first-price auction rule.

491. After the eventual migration to the Unified First Price Auction was completed in September 2019, all bidders—including AdX bidders, header bidders, and non-Google exchanges using Open Bidding—competed on the same first-price basis, with the highest bidder paying its bid.[971] This change eliminated the inefficiencies and confusions caused by differences in auction formats, and removed the so-called "last look" over header bidding. It reduced transaction costs both for both bidders (who no longer needed to bid

[971] Comms Doc, "Ad Manager Unified 1st Price Auction" (Sep. 27, 2019), GOOG-DOJ-09714662, at -663 ("After the transition is complete, all publisher traffic is on 1st auction").

differently in different exchanges) and publishers (who no longer needed to inflate header bids to set value CPMs for a second-price auction).

492. To ease the transition for its bidder customers to the new auction rules, Google Ads and DV360 introduced new programs to optimize their bids into the Unified First Price Auction. As discussed in Section IV.C.1.c, Alchemist optimized bids into the UFPA for Google Ads advertisers, while using threshold pricing to determine payments by advertisers.[972] The threshold prices made the Google Ads internal auction bidder-truthful for those advertisers. On DV360, Google determined bids into the UFPA on behalf of its advertisers (unless the advertiser opted out) using a bid optimization program similar to earlier versions of Poirot.[973] A winning DV360 advertiser would be charged its bid (plus any applicable platform fees).[974] This payment rule also incentivizes truthful reporting of campaign parameters so long as DV360 is trusted to determine the optimal bid shading factors on behalf of its advertisers. In combination, these programs maintained the simplicity of Google's buy-side tools, while allowing the bids from different demand-side platforms to be directly and simply compared.

493. Since the transition to the UFPA, Google has provided real-time bidders—including Authorized Buyers and Open Bidders—with historical auction information to allow

[972] Recall that this means that a Google Ads advertiser pays the lowest value it could have reported while still winning the impression. *See* Design Doc, "The Alchemist (AKA First Price Bernanke)" (Mar. 2019), GOOG-DOJ-14550102, at -103 to -104 (providing details on how the payment is calculated); Declaration of [redacted] (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 22.

[973] Declaration of [redacted] (Aug. 5, 2023), GOOG-AT-MDL-008842383, at ¶ 35 ("With the transition to a Unified First Price Auction, Google began providing minimum-bid-to-win data to buyers, and DV360 began to use that minimum-bid-to-win data to inform how Poirot would lower bids into AdX in order to optimize for expected advertiser surplus.").

[974] Or as otherwise agreed in the contract made between the advertiser and DV360, but this was the standard way advertisers were charged by DV360. *See* "ORDER FORM - DoubleClick Bid Manager Service" (Dec. 10, 2015), GOOG-DOJ-09457868, at -868 to -869.

non-Google buying tools to optimize their bids into the UFPA.[975] Among this information is the **minimum bid to win (MBTW)**: the smallest amount a bidder needs to bid to win an auction, holding all other bids in the auction fixed.[976] Google shares this MBTW data with *all* real-time bidders with which it shares a direct communication channel, including third-party DSPs and Open Bidding exchanges.[977]

494. In addition to sharing MBTW data with ad buyers, Google has also provided its publisher customers with auction data via Bid Data Transfer (BDT) files.[978] Early versions of these files were implemented prior to the UFPA, with the aim of providing publishers information about buyers' bids.[979] However, Google and its AdX customers had concerns about publishers' ability to access fields that allowed for buyers' losing bids to be matched to sensitive end-user information.[980] For example, [redacted] subjected Google to

[975] *See* Google, "2019 Google Ad Manager releases archive," Google Ad Manager Help (Jul. 1, 2019), https://support.google.com/admanager/answer/9197913?hl=en#expand&expand070119 ("Buyers currently have the ability to access feedback on prior bids they submitted into the auction […] As part of the transition to a first-price auction, buyers will gain access to a similar field (minimum_bid_to_win) that provides feedback on the minimum bid value necessary to have won a prior auction.").

[976] For the auction winner, this value is the larger of the runner-up bid and the floor price. For auction losers, the value is the larger of the highest bid and the floor price. *See* Presentation, "UPR/1P for Buyers gTech Resources" (Jun. 24, 2019), GOOG-DOJ-14039426, at -431.

[977] *See* Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 11 ("The Unified First-Price Auction sends the minimum-bid-to-win information to real-time bidders that directly bid into the Unified First-Price Auction (as long as their bid was not filtered from the auction), including all third-party exchanges that participate in Open Bidding and all bidders in Google's ad exchange.").

[978] *See* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 11.

[979] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 11 ("Prior to September 2019, Google allowed buyers to opt out of including information about their bids in BDT files."). *See also* Launch Details Spreadsheet, Launch 199259 (Aug. 29, 2023), GOOG-AT-MDL-009644182, at cells C2, C3 (noting launch date of 2017-9-5 for "Jedi Bid Data Transfer Beta").

[980] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13 ("The inclusion of information about a more complete set of bids raised two types of concerns within Google. First, there were concerns that a then-existing contract with an important Authorized Buyer prohibited Google from providing publishers with information about its bids in a format that could be joined with other Data Transfer report files. Second, there were privacy concerns that certain types of bid data could be tied to individual users if BDT files containing a more complete set of bids could be combined with other types of Data Transfer files.").

contractual restrictions that "disallow[] Google from sharing [redacted] losing bids tied to user identifiable information[.]"[981] Google addressed these privacy and contractual concerns by enabling buyers to opt out of sharing bid data with publishers, but that led to a significant portion of bids to be excluded from the datasets.[982] As Google later transitioned to the UFPA, it sought "to provide publishers more information about the new first-price auction by including a *complete* set of bids from all Authorized Buyers, Open Bidding buyers, Google Ads, and DV360."[983] With the "inten[tion] to provide publishers more information" in the move to the UFPA,[984] Google modified BDT files to prevent linking between losing bids and end-user data[985] while also "remov[ing] the

[981] Presentation, "1P Auction - Bid Data Transfer[:] Roll-out Plan" (Jul. 24, 2019), GOOG-DOJ-06882418, at -420 ("New Bid Data Transfer file, providing publishers with complete visibility into every single programmatic bid on every auction (all buyers on AdX + EB, remove buyer opt-out from bid sharing) ... but have since realized this is incompatible with existing contractual obligations to [redacted] has a contract clause that disallows Google from sharing [redacted] losing bids tied to user identifiable information[.]"); *see* Buyer Terms Agreement, "Google DoubleClick Ad Exchange Buyer Terms" (Apr. 8, 2014), GOOG-DOJ-15247796, at -796, -799 ("These Google DoubleClick Ad Exchange Buyer Terms ('Terms') are entered into by Google Inc. ('Google') and A9.com, Inc. ('Customer') [...] Customer Data will be deemed 'Confidential Information' of Customer. Notwithstanding the foregoing, (i) Google may disclose *non-User Identifiable* Customer Data to Partners if Google generally discloses similar data from other buyers to Partners in a substantially similar manner under the Program, (ii) Google may disclose *non-User Identifiable* Customer Data to report statistics about the Program if such data is in aggregate form and not identified or identifiable as Customer's or Advertiser's, and (iii) with respect to an auction for any impression for which Customer is the *winning bidder*, Google may disclose the User-Identifiable Customer Data related to Customer's bid request to the Partner on whose Publisher Property the Ad is served. If Google offers a new Program feature that discloses Customer Data to Partners, then within 30 days of a Customer request, Google will provide a demonstration or written description of the nature and extent of Customer Data being disclosed through that new Program feature.") (emphasis added). *See also* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13 ("First, there were concerns that a then-existing contract with an important Authorized Buyer prohibited Google from providing publishers with information about its bids in a format that could be joined with other Data Transfer report files.").

[982] *See* [redacted] GOOG-AT-MDL-C-000073682, at ¶ 11 ("For example, on August 31, 2019, approximately 70% of impressions and 50% of publisher revenue came from buyers that had opted out of having their information included in BDT files.").

[983] [redacted] GOOG-AT-MDL-C-000073682, at ¶ 12 (emphasis added).

[984] *See* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 12.

[985] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13 ("The inclusion of information about a more complete set of bids raised two types of concerns within Google. First, there were concerns that a then-existing contract with an important Authorized Buyer prohibited Google from providing publishers with information about its bids in a format that could be joined with other Data Transfer report files. Second, there were privacy concerns that certain types of bid data could be tied to individual users if BDT files containing a more complete set of bids could be combined with other types of Data Transfer files. Because of these

option for bidders to opt out of including information about their bids[.]"[986] These combined changes ensured that Google fulfilled its contractual commitments, addressed the privacy concerns of its buyers, and also shared a more complete set of bids with its publisher customers.[987]

D. Responding to Plaintiffs' and Their Experts' Allegations

1. Plaintiffs Mischaracterize the Open Bidding Program

495. Plaintiffs allege that Open Bidding was devised to "maintain its exchange monopoly and exclude competition from exchanges,"[988] but Plaintiffs' allegations mischaracterize or do not account for relevant details of the Open Bidding program.

496. *First,* the Plaintiffs claim that Open Bidding "requir[es] publishers to route their inventory through AdX, even if they do not want to."[989] But this is incorrect. After Open Bidding, publishers could continue calling exchanges via header bidding. While Open Bidding and header bidding could function alongside one another, publishers also had the flexibility to fully opt out of Google's Open Bidding services.[990] Non-Google exchanges

---

two concerns, at the same time in September 2019 that Google updated BDT files to include information about a more complete set of bidders, Google also changed the 'KeyPart' field in BDT files to be unique to those files and not joinable with other Data Transfer report files, and Google removed other fields that otherwise could have been used to probabilistically join BDT files with other Data Transfer report files.").

[986] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 12 ("When Google transitioned to a Unified First Price Auction in September 2019, Google removed the option for bidders to opt out of including information about their bids in BDT files.").

[987] *See* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13.

[988] Fourth Amended Complaint ¶ 367.

[989] Fourth Amended Complaint ¶ 370.

[990] *See* Happy Das, "Header Bidding vs Open Bidding – What Should You Know!," headerbidding.co Blog (Feb. 23, 2024), https://headerbidding.co/header-bidding-vs-open-bidding/ ("There are pros and cons to both Open Bidding and header bidding, so the right solution will ultimately depend on what you're looking for. If you can't seem to decide between the two, consider them both as options, as they can ultimately coexist within your ad server. Fortunately, you're never obligated to choose one over the other—make the call that works best for your setup, even

could also choose not to participate in Open Bidding, and publishers could still access those exchanges through header bidding. For these reasons, Professor Weinberg's "[c]omparison of header bidding to Exchange Bidding" is not relevant.[991] A publisher need not select between header bidding or Open Bidding. It can always opt to use header bidding instead of or alongside Open Bidding—whichever it finds more profitable.

497. *Second,* Plaintiffs argue that "[h]eader bidding lets each exchange access a cookie on the user's page, which permits an exchange to recapture some information about the user's identity," while "Exchange Bidding diminishes the ability of non-Google exchanges […] to identify users associated with publishers' heterogeneous inventory."[992] However, Google provided Open Bidding partners with its Cookie Matching Service that enabled exchanges to "match [...] their cookie[s]" for many impressions.[993] Additionally, publishers and exchanges could continue to use header bidding, leaving any existing implementation of end-user identification unchanged.

498. *Third*, Plaintiffs express concerns about the Open Bidding revenue share, but they fail to explain why a [redacted] revenue share is too much to charge for Google's Open Bidding

if it's called something different. You can always change course in the future if you find yourself wishing you had picked a different solution from the get-go."). *See also* Ad.Plus, "Header Bidding vs Open Bidding," Ad.Plus Blog (Jan. 6, 2023), https://blog.ad.plus/header-bidding-vs-open-bidding/ ("Google Open Bidding is often used in conjunction with header bidding, which allows multiple demand sources to bid on ad inventory simultaneously. It is intended to be a more efficient way for publishers to monetize their ad inventory, as they can potentially receive higher bids from a wider pool of buyers.").

[991] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at Section V.B.2.

[992] Fourth Amended Complaint ¶ 367.

[993] Google for Developers, "Cookie Matching," Real-time Bidding (accessed Jul. 17, 2024), https://developers.google.com/authorized-buyers/rtb/cookie-guide ("Open Bidding allows exchanges to use bidder initiated [...] and Google initiated [...] cookie matching workflows, to match a Google User ID with their cookie."). *See also* [redacted] (Redacted) (Sep. 29, 2023), GOOG-AT-MDL-C-000016753, at ¶ 16.

services.[994] For publishers, these services included reporting, payment processing, and integration with non-Google exchanges; for exchanges, they included real-time processing of the huge number of bids they submitted on each impression.[995] Moreover, other non-Google header bidding services charge fees. For example, Amazon's UAM deducts a 10% fee from bids into its header bidding service.[996] Internally, Google employees assessed that other exchange bidding tools charged revenue shares of between [redacted][997] Publishers who wished to avoid the Open Bidding revenue share could choose to use header bidding instead.

499. *Fourth*, Plaintiffs claim that, "in operating Exchange Bidding, Google maintained visibility into the bids submitted by rival exchanges and used that information to inform its own trade decisions" and "could win an auction for a publisher's inventory even over another exchange's higher bid."[998] These statements are incorrect in three ways.

i. Before the 2019 transition to a Unified First-Price Auction, neither Google Ads nor DV360 ever used bids from Open Bidding or header bidding to adjust the amount of a bid on the same impression.[999]

---

[994] Fourth Amended Complaint ¶ 369.

[995] Comms Doc, "Open Bidding on Ad Manager (fka Exchange Bidding)" (Aug. 2019), GOOG-DOJ-15389438, at -438 ("Easy to set up, view/analyze reports and unified payments […] Allows exchanges to respond to RTB call-outs […] Provides integrated reporting and billing for exchange bidding transactions won by 3rd party exchanges").

[996] *See* Amazon, "Unified Ad Marketplace," Amazon Publisher Services (accessed Sep. 12, 2023), https://aps.amazon.com/aps/unified-ad-marketplace/index.html ("UAM charges a 10% transaction fee from SSP and Amazon bid prices prior to conducting a first price auction.").

[997] Untitled Document (Mar. 9, 2021), GOOG-DOJ-AT-01502155, at -155 ("[redacted].").

[998] Fourth Amended Complaint ¶ 371.

[999] Even if the floor price faced by Google Ads or DV360 was set by a non-Google bidder (such as in header bidding or in the initial release of Open Bidding), neither Google Ads nor DV360 used it to adjust their bids. *See* [redacted] (May 01, 2024), GOOG-AT-MDL-C-000017969, at ¶¶ 2-3 ("Before the conclusion of an AdX auction

ii. Since March 2017—over one year before the general release of Open Bidding—AdX has not had a "last look" or "visibility into the bids" of Open Bidding exchanges,[1000] and the transition to the Unified First-Price Auction eliminated any so-called "last look" over bids from header bidding exchanges.

iii. Since the general launch of Open Bidding, the exchange with the highest bid received under Open Bidding (net of the appropriate revenue share) was awarded the impression.[1001]

2. Open Bidding Did Not "Foreclose" Competition from Header Bidding or Other Ad Exchanges

500. Plaintiffs claim that Google's introduction of Open Bidding "successfully forecloses competition from header bidding and in the exchange market,"[1002] but external data reported by eMarketer on the top 1,000 US publishers (reproduced as Figure 15 below) suggests that the proportion of all publishers using header bidding continued to grow after Open Bidding launched in 2018. This is consistent with Google's analysis of its Open Bidding program, which suggested that many publishers increasingly adopted

---

for an ad impression, Google Ad Manager has never provided Google Ads or DV360 with any bids that AdX or Open Bidding received from non-Google ad exchanges or non-Google buy-side tools in that auction. Before Google Ad Manager transitioned to a Unified First-Price Auction in 2019, Google Ads and DV360 never used a floor price in an auction for an ad impression to adjust the bid values that they would submit for that same ad impression.").

[1000] "2018 Sellside Launches Revenue Evaluation" (Jul. 19, 2019), GOOG-DOJ-13226855, at tab "Q1Q2 '17 Launch News," cells B60, R60 (noting launch date of 3/29/17 for "Remove Last Look for Demand Syndication Auction").

[1001] *See* Launch Document, "Launch Doc: Adx Auction in Post Revshare Space" (Jan. 31, 2017), GOOG-DOJ-AT-02425378, at -378 ("In the post-revshare auction, we pick a buyer that maximizes ability to pay after Google's revshare has been discounted."); Launch Details Spreadsheet, Launch 181133 (Aug. 26, 2023), GOOG-AT-MDL-009644157, at cell C2 (noting launch date of 2017-3-2).

[1002] Fourth Amended Complaint ¶ 372.

allocation methods incorporating more real-time bids after the introduction of Open Bidding (using either header bidding or Open Bidding or a combination of the two).[1003]

Figure 15: External data from eMarketer suggests that the adoption of header bidding by publishers was not negatively impacted by introduction of Open Bidding.[1004]

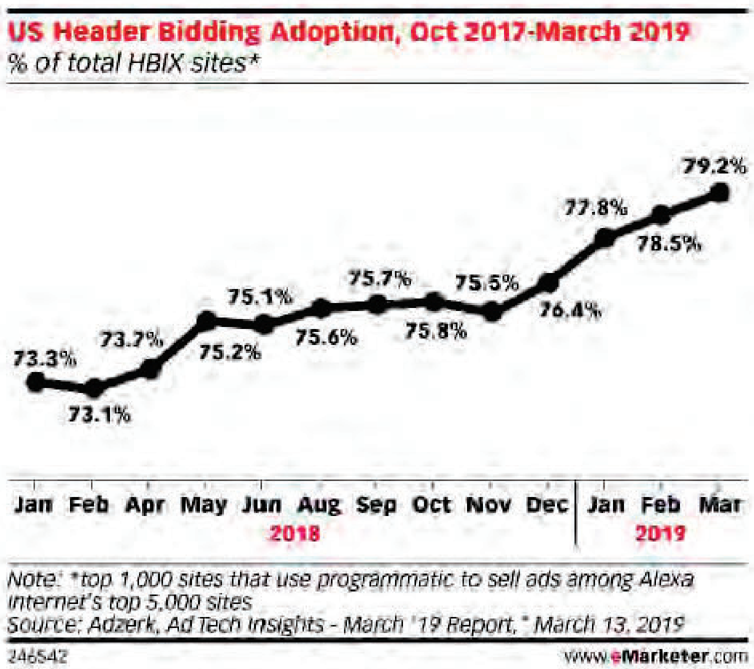

501. Data from Google's Header Bidding Monitor dataset also suggests that publishers continued utilizing header bidding after adopting Open Bidding. According to the dataset, as of March 2023, [redacted] North American publishers using Google Ad Manager used Open Bidding for some impressions. Of those [redacted] publishers

[1003] Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -775 ("Publishers are moving from a waterfall based approach to RTB (EB and Header Bidding)").

[1004] *See* Ross Benes, "Five Charts: The State of Header Bidding," eMarketer Insider Intelligence (May 30, 2019), https://www.insiderintelligence.com/content/five-charts-the-state-of-header-bidding.

using Open Bidding, [redacted]) sold impressions via channels that Google attributed to header bidding.[1005]

502. Open Bidding also made it *easier* for publishers to partner with competing exchanges, allowing exchanges like OpenX and Index Exchange to expand the number of impressions they could serve.[1006] Simplifying payment processing for publishers was seen as one of the key advantages of Open Bidding.[1007] The reduced costs of integrating and processing bids from the non-Google exchanges participating in Open Bidding encouraged more publishers to offer impressions to multiple exchanges, which benefited those exchanges.[1008]

[1005] This statistic is based on [redacted], *see* GOOG-AT-EDVA-DATA-000000006, grouped by month, with Open Bidding entries taken to be those with a transaction_type of "EBDA" or numbers 9-11, and header bidding entries taken to be those with an [redacted] *See* [redacted] Appendix A (Oct. 6, 2023), GOOG-AT-MDL-C-000012826, at -874. Publishers are grouped by their network_id and ranked according to the number of queries matched in the dataset. Code for this result can be found in code/hb_monitor_ob_freq.py in my supporting materials, and the output is saved in code/logs/hb_monitor_ob_freq.txt.

[1006] *See* Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -761 ("[redacted] […]").

[1007] *See* Sagar Bhatia, "Google Open Bidding (OB) Explained," AdSparc (Apr. 12, 2022), https://adsparc.com/google-open-bidding-explained/ ("Here are some advantages of Google Open Bidding over header bidding: Payment Management: One of the critical advantages of using OB over header bidding is payment management. For OB, payments are sent directly to publishers via Google, whereas for header bidding, payments are managed by vendors or the publishers themselves […]").

[1008] *See* Presentation, "Exchange Bidding Sell Side Update" (Jun. 14, 2018), GOOG-DOJ-11790760, at -775 (Table, "[redacted]").

3. Plaintiffs' and Their Experts' Incorrectly Assert that "Minimum Bid to Win" Data Disadvantages Header Bidders

a) Sharing Minimum Bid to Win Does Not "Recapture" a "Last Look Advantage"

503. Plaintiffs allege that Google and Open Bidding exchanges could use MBTW data "to adjust their future bidding strategy to continue trading ahead of exchanges returning bids through header bidding and underpaying for publishers' impressions"[1009] and that this data allowed Google to "recapture the advantages it had under Last Look."[1010] Plaintiffs' experts seem to agree, with Professor Pathak writing that, "although Google removed its Last Look advantage over Header Bidding from the unified auction, Google introduced a program with a similar effect by sharing 'Minimum Bid to Win' data with AdX and Exchange Bidding buyers and not Header Bidding buyers."[1011] These claims are incorrect.

504. *First,* there is no "advantage" to "recapture" under "Last Look."[1012] As I have explained in Section X, taking proper account of bidder and publisher incentives, the so-called "last

[1009] Fourth Amended Complaint ¶ 380 ("Specifically, in 2019, DFP began sharing sensitive pricing information derived from publishers' sensitive clearing auction records (which Google called "Minimum Bid to Win" data) with exchanges in Exchange Bidding. Google's AdX exchange and other exchanges in Exchange Bidding use this data to adjust their future bidding strategy to continue trading ahead of exchanges returning bids through header bidding and underpaying for publishers' impressions.").

[1010] Fourth Amended Complaint ¶ 381 ("Google compounded this Exchange Bidding advantage with a new secret bid optimization scheme that allowed Google to recapture the advantages it had under Last Look. The new scheme uses information about publishers' ad server user IDs and rival exchanges' bids to accurately predict the amount to bid, effectively permitting Google to re-engineer the ability of AdX and Google's ad buying tools to trade ahead of rivals exchanges in Exchange Bidding. As a Google planning document outlines: 'If we knew our competitor's bid exactly, we can simply bid a cent above that[.] But we don't have this information before the auction, so we need to predict [the] competitor's bid.'").

[1011] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 199 ("For instance, although Google removed its Last Look advantage over Header Bidding from the unified auction, Google introduced a program with a similar effect by sharing 'Minimum Bid to Win' data with AdX and Exchange Bidding buyers and not Header Bidding buyers. Minimum Bid to Win data effectively recreates Last Look, because AdX and Exchange Bidding buyers can use the information from Minimum Bid to Win on their next auctions, while Header Bidding buyers cannot.").

[1012] Fourth Amended Complaint ¶ 381.

look" did not inherently advantage AdX and other Open Bidding exchanges. Moreover, any suggestion that AdX or Open Bidding buyers could "peek" at a competitors' bids before submitting their bids at the time of the UFPA is wrong.[1013] To the contrary, the MBTW was shared *after* the auction's conclusion.[1014]

505. *Second*, any suggestion that using MBTW information is nefarious and somehow enabled AdX or other Open Bidding buyers to exactly predict header bids and "displace" their trades by "a penny" is wrong. In first-price auctions, bidders' competitors vary their participation and bids from impression to impression—even for impressions from the same end user—o MBTW information does not enable a bidder to "know [a] competitor's bid exactly"[1015] and "displace" them "by a penny."[1016] It is widely understood that optimal bidding in a first-price auction requires bidders to make

[1013] Fourth Amended Complaint ¶¶ 378-81 ("Starting with the official launch of Exchange Bidding in June of 2017, Google sought to lure exchanges away from header bidding by sharing its Last Look advantage with other exchanges participating in Exchange Bidding. These exchanges could now also peek at header bidding net bids and displace their trades by a penny [...] Several years later, [...] DFP began sharing sensitive pricing information derived from publishers' sensitive clearing auction records (which Google called 'Minimum Bid to Win' data) with exchanges in Exchange Bidding. [...] Google compounded this Exchange Bidding advantage with a new secret bid optimization scheme that allowed Google to recapture the advantages it had under Last Look. The new scheme uses information [...] to accurately predict the amount to bid, effectively permitting Google to re-engineer the ability of AdX and Google's ad buying tools to trade ahead of rivals exchanges in Exchange Bidding.").

[1014] Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 10 ("The minimum-bid-to-win is information that the Unified First-Price Auction sends after a given auction has ended.").

[1015] Fourth Amended Complaint ¶ 381 ("The new scheme uses information about publishers' ad server user IDs and rival exchanges' bids to accurately predict the amount to bid, effectively permitting Google to re-engineer the ability of AdX and Google's ad buying tools to trade ahead of rivals exchanges in Exchange Bidding. As a Google planning document outlines: 'If we knew our competitor's bid exactly, we can simply bid a cent above that[.] But we don't have this information before the auction, so we need to predict [the] competitor's bid.'").

[1016] Fourth Amended Complaint ¶¶ 378-81 ("Starting with the official launch of Exchange Bidding in June of 2017, Google sought to lure exchanges away from header bidding by sharing its Last Look advantage with other exchanges participating in Exchange Bidding. These exchanges could now also peek at header bidding net bids and displace their trades by a penny. Several years later, [...] DFP began sharing sensitive pricing information derived from publishers' sensitive clearing auction records (which Google called "Minimum Bid to Win" data) with exchanges in Exchange Bidding. [...] Google compounded this Exchange Bidding advantage with a new secret bid optimization scheme that allowed Google to recapture the advantages it had under Last Look. The new scheme uses information [...] to accurately predict the amount to bid, effectively permitting Google to re-engineer the ability of AdX and Google's ad buying tools to trade ahead of rivals exchanges in Exchange Bidding.").

probabilistic assessments of competitors' likely bids[1017] to determine their own bids, and Google made this easier for its customers by providing feedback from past auctions.

b) Plaintiffs' and Their Experts' Arguments are Ahistorical and Fail to Account for Participants' Alternatives

506. Plaintiffs' and their experts' arguments suggest that not "sharing" MBTW data with "Header Bidding buyers" disadvantaged them.[1018] However, this argument overlooks the technological limitations of "sharing" MBTW data with header bidding exchanges, ignores the alternative methods available for header bidding buyers to obtain this data, and overstates the unique importance of MBTW data.

507. As a publisher-configured technology, header bidding did not provide GAM with a direct communication channel to header bidding exchanges. As noted by a Google engineer, "[h]istorically, there was no basis on which this information could be sent by Ad Manager to header bidders because there is no server-to-server connection between Ad Manager and such bidders[.]"[1019] Instead, GAM enabled publishers to obtain a copy of MBTW information, which they could share directly with header bidding exchanges.[1020]

[1017] *See, e.g.*, Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37, at 21 ("In the top-price method of negotiation, as in the Dutch auction, bidders, in order to maximize their expectation of profit, must concern themselves not only with their own appraisal of the article but also with their estimate of the value that others will place on it and their expectation of the bidding strategy that others will follow. This involves a considerable amount of appraisal of the market situation as a whole[.]").

[1018] *See, e.g.*, Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 199 ("Google introduced a program [... that] shar[ed] 'Minimum Bid to Win' data with AdX and Exchange Bidding buyers and not Header Bidding buyers.").

[1019] Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 12.

[1020] Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 12 ("Historically, there was no basis on which this information could be sent by Ad Manager to header bidders because there is no server-to-server connection between Ad Manager and such bidders and Ad Manager does not know the identify of these exchanges (although publishers could pass that information along to header bidders)."); Internal Briefing, "Briefing: News Corp Unified Pricing Floors Discussion" (Apr. 18, 2019), GOOG-DOJ-AT-00045716, at -721 ("If you use HB, you

508. There were also additional methods for header bidding buyers to obtain MBTW data:

a. *Open Bidding*. Header bidding exchanges can acquire MBTW information by submitting bids into Open Bidding.[1021] If the exchange wishes to transact impressions only through header bidding, it can do that and still receive MBTW information by submitting lower bids through Open Bidding.

b. *Multi-homing DSPs*. Plaintiffs incorrectly presume that *exchanges* submit bids, rather than *DSPs*. When DSPs submit bids, they frequently multi-home, submitting bids into numerous exchanges on the same impression opportunity.[1022] A DSP that bids into both a header bidding exchange and AdX receives the MBTW from GAM.[1023] The DSP is not required to purchase any additional impressions through AdX to acquire this information; it can bid lower into AdX as a means to receive MBTW information.

c. *Google Ad Manager*. In June 2022, Google introduced a feature enabling "header bidders who participate in the Ad Manager auction for web inventory" to receive MBTW information directly from GAM.[1024]

can use Data Transfer [files], where you can see all of the bids and the winning bid, and you can share this with anyone you like.").

[1021] *See* Comms Doc, "Header bidding in yield groups" (May 3, 2022), GOOG-AT-MDL-006196134, at -136 ("Currently [minimum bid to win] information is shared by Ad Manager - after an auction is run - with any Authorized Buyers and Open Bidders that submitted a bid in that auction.").

[1022] *See, e.g.*, [redacted]

[1023] Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 11 ("The Unified First-Price Auction sends the minimum-bid-to-win information to real-time bidders that directly bid into the Unified First-Price Auction (as long as their bid was not filtered from the auction), including all third-party exchanges that participate in Open Bidding and all bidders in Google's ad exchange.").

[1024] Declaration of N. Korula (May 2, 2024), GOOG-AT-MDL-C-000017971, at ¶ 12. *See also* Comms Doc, "Header bidding in yield groups" (May 3, 2022), GOOG-AT-MDL-006196134, at -136 ("This feature allows header bidding

509. In addition to overlooking all the ways in which header bidding buyers could access MBTW data, Plaintiffs and their experts significantly overstate the uniqueness of that data. Buyers without access to MBTW data could use other forms of feedback to effectively optimize their bids into first-price auctions. For example, such a buyer could experiment to learn optimal bids for the UFPA, as DV360 did for non-Google exchanges under Poirot (see Section VII).

4. Plaintiffs' and Their Experts' Bid Data "Redaction" Allegations Are Historically Incomplete and Fail to Account for Publishers' Alternatives

510. Plaintiffs and their experts claim that, in 2018, Google began redacting information from data files that it provided publishers, "crippl[ing] publishers' ability to measure the success of rival exchanges in header bidding[.]"[1025]

511. *First,* Plaintiffs' and their experts' description is historically incomplete and omits various contractual and privacy considerations. As I summarize in Section XIII.C.3, Google "intended to provide publishers [with] more information about the new first-price auction by including a complete set of bids from all Authorized Buyers, Open Bidding buyers, Google Ads, and DV360"[1026] as "approximately [redacted] and [redacted]

partners to receive [minimum bid to win] information, representing the minimum bid value necessary to have won an auction.").

[1025] Fourth Amended Complaint, Section VII.D.3.iii ("Google cripples publishers' ability to measure the success of rival exchanges in header bidding (2018 to present)."). *See also* Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 150 ("The redaction of the DT protected AdX against the threat of Header Bidding because it removed publishers' ability to measure Header Bidding results and effectively target users."); Expert Report of J. Gans (Jun. 7, 2024), at ¶¶ 683, 685 ("In short, by changing the way that KeyPart was generated (or 're-encoding' it), Google removed the ability for publishers to match DT files together [...] By redacting these fields, Google prevented publishers from knowing relevant identifying information about auction winners.").

[1026] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 12.

publisher revenue came from buyers that had opted out" at the time.[1027] Before Google could provide its publisher customers with "a more complete set of bids," it needed to resolve additional privacy and contractual concerns.[1028] Ultimately, modifying certain fields in the BDT files did just that, enabling Google to securely share bid information of more buyers.[1029]

512. Professor Gans disregards the historical context of these changes, undermining his conclusion that "Google's claim that it redacted data based on privacy concerns is pretextual" and that the "sole intent [...] was to remove the ability of publishers to join these files together and gain insights about their businesses and performance across various exchanges."[1030] Instead, he attempts to support his conclusion by quoting a single document[1031] that states, "We want to prevent a publisher being able to determine 'these

[1027] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 11.

[1028] *See* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13.

[1029] *See* [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 13 ("The inclusion of information about a more complete set of bids raised two types of concerns within Google. First, there were concerns that a then-existing contract with an important Authorized Buyer prohibited Google from providing publishers with information about its bids in a format that could be joined with other Data Transfer report files. Second, there were privacy concerns that certain types of bid data could be tied to individual users if BDT files containing a more complete set of bids could be combined with other types of Data Transfer files. Because of these two concerns, at the same time in September 2019 that Google updated BDT files to include information about a more complete set of bidders, Google also changed the 'KeyPart' field in BDT files to be unique to those files and not joinable with other Data Transfer report files, and Google removed other fields that otherwise could have been used to probabilistically join BDT files with other Data Transfer report files.").

[1030] *See* Expert Report of [redacted] (Jun. 7, 2024), at Section VII.D.2 ("Google's claim that it redacted data based on privacy concerns is pretextual"), ¶ 689 ("Google's sole intent in making the changes to the DT files was to remove the ability of publishers to join these files together and gain insights about their businesses and performance across various exchanges.").

[1031] *See* Expert Report of [redacted] (Jun. 7, 2024), at ¶ 690 ("In a 2019 document, Google employees discuss the changes made to DT files. To the question, 'Why do we redact and roun[d] data?' an employee explains, 'We want to prevent a publis[her to be] able to determine "these advertisers were will[ing to pay] this much for that user's impression.' Google's intent is, hence, to limit the amount of information publishers receive to prevent them from establishing advertisers' willingness to pay across competing exchanges.") (quoting GOOG-NE-04599495, at -495, GOOG-DOJ-04602757, at -757).

advertisers were willing to pay this much for that user's impression.'"[1032] However, rather than providing support for his conclusion, that quote aligns with the historical context I have provided regarding Google's concerns for user privacy, its contractual commitments, and its desires to share the bids of more buyers with publishers. Indeed, the same document notes that Google's goal was to "to give publishers full transparency into the auction, while still protecting user privacy and the disclosure of proprietary data or user data that may be reflected in buyer's bids (for example, user interactions on their websites)."[1033]

513. *Second,* Plaintiffs' and their experts' allegations overstate the significance of the "redactions" to the BDT files.

a. After the "redactions," publishers could still utilize BDT files to infer whether bids from AdX or Open Bidding had lost to a competing line item, including a header bidding exchange. BDT files contain a field called "BidRejectionReason" which is set to "Outbid" if the bid "lost to another candidate in the auction or [a] competing ad server line item."[1034] Since header bidding was configured using ad server line items, a high-bid marked as "Outbid" could suggest that the

[1032] "1P Bid Data Transfer" (Aug. 9, 2019), GOOG-NE-04599495, at -495, GOOG-DOJ-04602757, at -757 ("We want to prevent a publisher being able to determine 'these advertisers were willing to pay this much for that user's impression.'").

[1033] "1P Bid Data Transfer" (Aug. 9, 2019), GOOG-NE-04599495, at -495, GOOG-DOJ-04602757, at -757 ("We want to give publishers full transparency into the auction, while still protecting user privacy and the disclosure of proprietary data or user data that may be reflected in buyer's bids (for example, user interactions on their websites).").

[1034] Google, "Bids data in Ad Manager Data Transfer (Beta)," Google Ad Manager Help (accessed Jul. 24, 2024), https://support.google.com/admanager/answer/7357436?hl=en ("BidRejectionReason[:] Reason the bid lost or did not participate in the auction. Possible values include: [...] 'Outbid': The bid lost to another candidate in the auction or competing ad server line item.").

impression was sold to a demand source represented by a line item with a high value CPM for the impression, such as a header bidding exchange.

b. Furthermore, publishers wishing to "evaluate the performance of exchanges in header bidding" had alternatives to acquire relevant information.[1035] For example, publishers could use A/B tests or experiments to evaluate how key performance indicators would change in the presence of certain header bidding exchanges.

514. BDT files have been a "beta" feature even as recently as 2024,[1036] and adoption by publishers has been [redacted] with [redacted] publishers subscribing to receive BDT files as of August 2019.[1037]

5. Professor Weinberg's "Last Look" Allegations Misrepresent the Auction Process

515. Professor Weinberg argues that "AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding"[1038] and that Open Bidding "creates two tiers" with header bidding exchanges in the "lowest tier" and AdX and Open Bidding exchanges in the "top tier."[1039] However, Professor Weinberg

[1035] Fourth Amended Complaint ¶ 387.

[1036] Google, "Bids data in Ad Manager Data Transfer (Beta)," Google Ad Manager Help (accessed Jul. 24, 2024), https://support.google.com/admanager/answer/7357436?hl=en ("This feature is in Beta").

[1037] [redacted] (Jun. 28, 2024), GOOG-AT-MDL-C-000073682, at ¶ 17 ("[redacted].").

[1038] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 161 ("AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding, but do not have a Last Look advantage over each other, placing them in the top tier.").

[1039] Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 161 ("Therefore, one interpretation of Exchange Bidding is that it creates two tiers: Exchanges that participate in header bidding and exchanges that participate in Exchange Bidding together with AdX. Exchanges that participate in header bidding submit bids without seeing others' bids, and therefore have no Last Look advantage over anyone, and are vulnerable to AdX's and Exchange Bidding's Last Look advantage (placing them in the lowest tier). AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding, but do not have a Last Look advantage over each other, placing them in the top tier.").

understates the possibility that a publisher could boost a header bid before it serves as the floor price in the ad auction to reduce or even reverse any purported "advantage."

516. As I discuss in Section X of this report, a publisher has incentives to increase ("boost" or "inflate") the header bids it submits in Open Bidding. When it does so, it increases the probability that the header bidding bid will win the impression, which makes it harder for a bidder on AdX or an Open Bidding exchange to win the impression. Depending on the extent of this bid inflation, there can be no "advantage" for AdX and Open Bidding exchanges. In fact, those exchanges could even be at a disadvantage. Despite acknowledging the possibility of header bid inflation, Professor Weinberg claims that "the impact [of a sophisticated publisher drastically inflating AdX's reserve specifically because of the Last Look advantage] is less clear-cut and would require a complicated analysis."[1040] But, the "impact" of boosting header bids *is* clear-cut. As I show in the technical notes in Section XV.F.1, in the example discussed by Professor Weinberg in his analysis of Open Bidding, a publisher that inflates header bids can achieve the same revenue as in a unified auction, while the bidders with the same distributions of values attain the same win rates and pay the same average prices. Consequently, there are no inherent tiers or "Last Look" advantages built into the Open Bidding design.

517. Professor Weinberg offers no evidence that publishers did not respond to their very clear incentives to boost header bids and no explanation as to why publishers would not

[1040] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 225 ("If a sophisticated publisher mildly inflates AdX's reserve specifically because of the Last Look advantage, or if advertisers bid similarly in these two cases, the same conclusions still qualitatively hold. If a sophisticated publisher drastically inflates AdX's reserve specifically because of the Last Look advantage, or advertisers drastically change their bids specifically due to AdX's Last Look advantage, the impact is less clear-cut and would require a complicated analysis weighing the benefits of Last Look versus the impact of an increased reserve and distinct bids.").

respond in that way. As I discuss in Section II.B.3 and Section X.D.2, there is evidence that publishers *did* boost header bids and there is good reason to expect that they would, given the importance of optimization for publisher returns and the simplicity of the logic involved.

518. Because AdX operated a second-price auction at the time, even if a publisher acts contrary to its incentives and does *not* boost header bids, it can still be better off under the so-called "last look" than under a scenario in which header bidders, AdX, and Open Bidding exchanges submit offers in a first-price auction without any "last look." If AdX submitted its clearing price without any "last look," the winning header bid would only need to exceed AdX's second-highest bid to win the impression. However, with "last look," the header bid must beat *both* the second-highest bid and the *highest* bid in AdX. This comparison would tend to encourage the header bidder to *increase* its bids in the auction with "last look" compared to the auction without, and that would increase publisher revenue. Indeed, in Section XV.G.2 of the Technical Notes, I demonstrate how the publisher revenues under "last look" are higher in a simple extension of Professor Weinberg's example, where the extension is that the exchange with "last look" receives bids from two bidders instead of just one.

XIV. UNIFIED PRICING RULES: PROTECTING ADVERTISERS FROM PRICE-FISHING BY PUBLISHERS USING EXCHANGE-DISCRIMINATORY FLOOR PRICES

A. Overview

519. In 2019, at the same time as its transition to a Unified First Price Auction (UFPA), Google introduced Unified Pricing Rules (UPR), which enabled publishers to use a single interface in Google Ad Manager to configure and manage floor prices that apply to all exchanges and demand sources.[1041] These floor prices could vary by properties of the impression and advertiser, but not by the identity of the exchange or the buying tool used by the advertiser.[1042]

520. UPR ensured that advertisers faced the same floor prices on all bidding channels in the UFPA. Before Google introduced the UFPA, publishers using GAM could improve both efficiency and revenue by setting different floor prices for bidders or demand sources

[1041] *See* Sam Cox, "Simplifying programmatic: first price auctions for Google Ad Manager," Google Ad Manager Blog (Mar. 6, 2019), https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/ ("[I]n the coming months we'll start to transition publisher inventory to a unified first price auction for Google Ad Manager."); Jason Bigler, "An update on first price auctions for Google Ad Manager," Google Ad Manager Blog (May 10, 2019), http://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/ ("In addition to impacting how publishers are using floor price rules, changing to a first price auction in Ad Manager requires a change in how our rules function. […] That's why we released a new feature to all publishers globally, called unified pricing rules.").

[1042] *See* Comms Doc, "Ad Manager Unified 1st Price Auction" (Sep. 27, 2019), GOOG-DOJ-09714662, at -665 ("Unified Pricing rules will not support the following functionalities that were present in Open Auction pricing rules: Buyer-specific floors: ability to set different floors for different buyers/bidders for a given inventory targeting […] publishers will still be able to: Set per-advertiser floors in Unified Pricing rules"); Google, "Unified pricing rules," Google Ad Manager Help (Jun. 25, 2024), https://support.google.com/admanager/answer/9298008 ("Advertiser- and brand-specific pricing can be configured in unified pricing rules. They don't apply to remnant line items. Per-buyer and per-bidder pricing are not available."); Jason Bigler, "An update on first price auctions for Google Ad Manager," Google Ad Manager Blog (May 10, 2019), https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/ ("To maintain a fair and transparent auction, these rules will be applied to all partners equally, and cannot be set for individual buying platforms.").

depending on the order in which they were called. This important justification of setting different floor prices for different exchanges was eliminated in the UFPA, as discussed by Google engineers.[1043]

521. UPR benefited Google's buyer-customers. Although Google's Open Bidding had addressed some of the flaws of header bidding (as described in Section XIII), advertisers still faced the risk of self-competition, in which an advertiser partners with multiple DSPs or bids into multiple exchanges when those are competing for the same impression. In turn, publishers could exploit this advertiser multi-homing through a tactic known as **price-fishing**: by setting different floor prices for different exchanges, a publisher could increase its revenue at the expense of such advertisers. Internal documents suggest Google was concerned about the possibility of price-fishing after the transition to the UFPA,[1044] and UPR protected advertisers from such tactics. Facing a common floor price also simplified the bidding process for advertisers and DSPs, reducing the costs of evaluating the same impression that might be offered many times at different floor prices

[1043] *See* Email from [redacted], "Fwd: First-price & Removing pricing knobs" (May 11, 2019), GOOG-DOJ-06732979, at -980 ("We know that these knobs become less useful/impactful in the first-price auction. So even if we keep them they will become ineffective and the usage should trend down in the coming months.").

[1044] *See, e.g.*, Email from [redacted]., "Fwd: 1st Price Changes" (Jun. 10, 2019), GOOG-DOJ-12948968, at -969 ("Having granular control also incentivizes pubs to call the same demand source multiple times throughdifferent channels with different floor prices (eg. DBM is called through AdX, EB and HB with different floor
prices), to effectively fish for the highest price.").

and formulating bid responses. Other online display advertising intermediaries have implemented similar rules requiring floor prices to be uniform across bidders.[1045]

522. UPR also benefited non-price-fishing publishers in two ways. *First*, it simplified the publisher's process of setting floor prices applying to both Google and non-Google buyers using Open Bidding. Instead of setting floor prices separately using each exchange's interface, a publisher needs to set just one UPR floor in GAM. *Second*, it mitigated an externality that would reduce the revenues of all publishers, arising when a multi-homing advertiser faced with price-fishing publishers finds it difficult to coordinate its many bids and is incentivized to reduce its bids or not bid at all on some impressions.

523. Plaintiffs and their experts fail to acknowledge the benefits of UPR to advertisers and publishers. A few months after the introduction of the UFPA and UPR, Google found that both total publisher revenues and the share of impressions allocated to non-Google exchanges increased.[1046] Post-launch experiments found that the combined effect of

---

[1045] *See, e.g.*, Meta, "Code of Conduct," Meta Audience Network (accessed Dec. 2, 2023), https://www.facebook.com/audiencenetwork/partner-program/code-of-conduct ("When a reserve price is applied as part of the auction, it should apply identically to all demand sources."). *See also* [redacted]

[1046] *See* Jason Bigler, "Rolling out first price auctions to Google Ad Manager partners," Google Ad Manager Blog (Sep. 5, 2019), https://blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/ ("Over the last few months, we've been testing the performance of this change and the results show that on average, first price auctions have a neutral to positive impact on a publisher's total revenue—revenue from all their advertising sources—when compared to second price auctions. In addition, we found evidence that first price auctions have created a more competitive market, resulting in third parties (Demand Side Platforms and Ad Networks outside of Google) and indirect line items (like those from Header Bidding implementations) winning an increased share of impressions.").

UFPA and UPR was [redacted] on publisher revenues and [redacted] for non-Google DSPs buying via AdX.[1047]

524. Instead, Plaintiffs and their experts claim that "UPR reduced the ability of publishers to maximize revenue by setting different reserve prices for distinct demand sources and limited their ability to ensure high-quality advertisements."[1048] But differential floor prices for different advertisers are still allowed under UPR and may be used to manage yield and filter lower quality ads.[1049] In addition, AdX also provides publishers several ways in which to control the quality and type of ads that can be displayed.[1050] Finally, publishers can and do use a *more effective* tool than differential floor prices to preference different exchanges, namely, **post-auction discounts**, which are contractual agreements made between publishers and exchanges or ad buyers to rebate a portion of winning bids.[1051]

[1047] *See* Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -772 ("[N]eutral publisher payment overall"); Presentation, "1PA Impact Post Launch" (Nov. 27, 2019), GOOG-DOJ-14566285, at -288 ("One month after launch, holdback experiment comparing 1PA to 2PA shows slightly lower revenues in general, but remains close to neutral across the Sell side"), -293 ("Top 3P Buyers: Positive impact due to lack of adjustment?").

[1048] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 19. *See also* Fourth Amended Complaint ¶ 453 ("Setting higher price floors for AdX and Google's buying tools permitted publishers to combat (but not solve) the problem of adverse selection caused by Google, thereby encouraging exchange and buyer participation (including those engaged in header bidding) and increasing overall yield. [...] Publishers also set high floors for Google's exchange and buying tools to diversify the sources of demand for their inventory. [...] Publishers also set higher price floors for AdX and Google's buying tools to improve the quality of the ads returned to their site and displayed to consumers.").

[1049] *See* Google, "Unified pricing rules," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/9298008 ("Advertiser- and brand-specific pricing can be configured in unified pricing rules.").

[1050] *See* Google, "Block sensitive categories," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/2541069?hl=en&sjid=14902226251943448609-EU ("You can block groups of ads that are considered 'sensitive' due to the nature of the business or ad […] Our system classifies ads automatically, and we don't rely on advertiser-provided categorization.").

[1051] *See* Seb Joseph, "WTF are post-auction discounts?," Digiday (Apr. 2 2020), https://digiday.com/media/post-auction-discounts/ ("Post-auction discounts happen when a publisher agrees to [t]ake money off of winning bids for certain impressions from an agency. So if an agency wins an impression in an auction at one price, the actual fee they pay will be lower once the discount kicks in."). Post-auction discounts were in use

525. Plaintiff and their experts claim that "[l]imiting the number of allowed reserve prices to 200 limits the publisher's ability to maximize revenues."[1052] This claim is incorrect. The limit is on "rules" and each rule allows for up to 50 distinct advertiser-specific reserve prices so the total number of allowed reserve prices is 10,000.[1053] In addition, Plaintiffs' experts do not present any evidence that this limit had any significant effect on publisher revenues.

B. Self-Competition and the Problem of Varying Floor Prices Across Exchanges

526. As explained in Section XIII, Open Bidding improved on header bidding by streamlining payments, simplifying configuration, and reducing computational burden on end users' browsers. However, advertisers still faced the risk of **self-competition**: that is, the possibility of competing against themselves for the same impression.

527. There were at least two ways in which an advertiser might end up competing against itself for the same impression.

528. *First*, by partnering with a *single* DSP that bid in multiple exchanges. Because, as acknowledged by Plaintiffs,[1054] DSPs often bid for inventory on multiple exchanges, header bidding could lead to advertisers bidding multiple times for the same impression.

---

by publishers using GAM from at least before 2020. *See* "Agency Programmatic Buying Models and Discounts" (Jun. 24, 2020), GOOG-DOJ-AT-00173317, at -318 ("Post-Auction Price Reduction is live with Rubicon, Index & OpenX.").

[1052] *See* Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 171.

[1053] *See* Google, "Unified First-Price Auction - Best practices," Google Ad Manager (accessed Jun. 15, 2024), https://services.google.com/fh/files/misc/unified_first-price_auction_best_practices.pdf ("You can currently specify up to 50 advertisers per pricing rule.").

[1054] Fourth Amended Complaint ¶ 71 ("Large advertisers do this with ad buying tools called demand-side platforms (commonly known as "DSPs"), which they use to optimize their spend across multiple exchanges and/or networks. Small advertisers, on the other hand, optimize and effectuate their purchases using pared-down analogues of DSPs.").

For example, if an advertiser partnered with a DSP that bid in the AdX auction and in header bidding auctions, the advertiser could find itself competing for the same impression through AdX and header bidding.[1055]

529. *Second*, by partnering with *multiple* DSPs who bid in the same exchange or auction. For instance, an advertiser might use The Trade Desk and Adobe, with each of those DSPs submitting bids on the advertiser's behalf in AdX as well as other exchanges. In situations like that, the same advertiser could end up competing against itself.

530. By itself, self-competition was not necessarily a problem for advertisers. Consider an advertiser who receives two bid requests for an impression—one each from Exchange A and Exchange B—with the same floor price of $1.50 in a unified first-price auction. The bid requests might be for either the *same* impression or *different* impressions that happen to have the same floor price, but the advertiser's bidding strategy is the same regardless: it bids as it would in a standard first-price auction (as analyzed in Section III.C.3.b) with a floor price of $1.50. For example, if the advertiser estimates that there are no other advertisers competing for the impression, then the advertiser would optimally place a bid in each exchange that *just* beats the floor price of $1.50 in either case, which is just the same as if there were no self-competition.

531. But self-competition *in combination* with varying floor prices across exchanges could harm advertisers if they (or their DSPs) fail to identify when their bid requests from

[1055] Presentation, "Optimal AdX in DFP setup: Best practices, and how to traffic RTA/RTP (header bidding) line items" (Sep. 24, 2015), GOOG-TEX-00000001, at -004 ("[H]eader bidding can make buyers bid against themselves running 2 auctions for every impression."); Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 41; [redacted]

different exchanges are competing for the same impression. The effect is similar to that of multi-calling, which was discussed in Section IV. To illustrate, consider the same example as before, but now suppose that the floor prices are \$1.50 in Exchange A and \$2.50 in Exchange B. *If* the advertiser can identify that both bid requests are for the *same* impression, then it would optimally bid \$1.50 in *both* Exchange A and Exchange B, winning the impression in the unified first-price auction (through Exchange A) at the lower price of \$1.50. *But if* the advertiser incorrectly thinks that the bid requests are for *different* impressions (when they are actually for the same impression), then it might bid \$1.50 in Exchange A and \$2.50 in Exchange B. In this case, the advertiser still wins the impression in the unified first-price auction (through Exchange B) but now pays the higher price of \$2.50: it is harmed by self-competition.

532. Avoiding self-competition before the introduction of UPR thus required DSPs to assess *many* similar or identical bid requests. One industry source reported that header bidding publishers might "send 18 identical bid requests for the same piece of inventory […] Currently, DSPs often begrudgingly evaluate them all, and find it hard to tell if the impression is exactly the same."[1056] Advertisers "struggle[d]" to de-duplicate impressions

[1056] Sarah Sluis, "The Trade Desk Suppresses Bid Duplication Amid COVID-19 Traffic Surge," AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge/ ("Publishers that slot the same six exchanges into multiple header-bidding auctions, such as Prebid, Google open bidding and Amazon Transparent Ad Marketplace, send 18 identical bid requests for the same piece of inventory to The Trade Desk. Currently, DSPs often begrudgingly evaluate them all, and find it hard to tell if the impression is exactly the same.").

and coordinate bids made through different channels.[1057] A Vice President of The Trade Desk commented that "[m]ore problematic is that each identical impression is represented differently to buyers specifically regarding floors. If floors represent a minimum price, why does an identical impression have 26 different floors?"[1058] Google employees were concerned that the possibility of self-competition could lead to an "[e]ventual loss of advertiser trust in [real-time bidding] auctions,"[1059] and that exchange-discriminatory reserves were a way for publishers to "game[]" the system.[1060]

533. Even though varying price floors across exchanges could harm advertisers, individual publishers had an incentive to do so. By engaging in price-fishing tactics, publishers could benefit at the expense of advertisers who participated on multiple exchanges. To illustrate the possible harms of such tactics, suppose that an advertiser's value for an impression could be either $2.20 or $1.80. In a first-price auction, the advertiser's optimal

[1057] *See* Presentation, "Unified 1st Price Auction" (Mar. 4, 2019), GOOG-DOJ-06525908, at -915 ("Today, buyers struggle to optimize when bidding across different channels due to lack of symmetry: different auction rules and different floor prices can apply for the same impression."). *See also* Sarah Sluis, "Header Bidding Unleashed a Huge Infrastructure Problem and Ad Tech Will Either Sink or Swim,"AdExchanger (Apr. 24, 2017), https://www.adexchanger.com/platforms/header-bidding-unleashed-huge-infrastructure-problem-ad-tech-will-either-sink-swim/ ("Though DSPs are seeing a flood of new impressions, many of them are selling the same thing. A publisher with three header bidding partners will have three exchanges sell its inventory, tripling the amount of impressions a DSP must evaluate and tripling the listening cost. […] [M]any DSPs are devoting resources to deduplicating impressions to avoid spending millions in server fees. […] But not all DSPs have the resources to smartly filter out low-value impressions.").

[1058] Will Doherty, "The door in the floors: Transparency, price discovery, and market efficiency in programmatic," The Current (Nov. 8, 2023), https://www.thecurrent.com/will-doherty-transparency-programmatic-data ("The average premium publisher works with approximately 26 different SSPs and exchanges. Our buyers now see every impression at least 26 times, and sometimes many more. More problematic is that each identical impression is represented differently to buyers specifically regarding floors. If floors represent a minimum price, why does an identical impression have 26 different floors?").

[1059] Presentation, "Optimal AdX in DFP setup: Best practices, and how to traffic RTA/RTP (header bidding) line items" (Sep. 24, 2015), GOOG-TEX-00000001, at -004.

[1060] *See* Email from [redacted], "Re: article - trial by fire, how 6 header bidders perform" (Jan. 11, 2016), GOOG-DOJ-13364449, at -449 [redacted] pubs can still play soft floor games even with HB or Jedi. They could still put different floors on different exchanges, calling them in parallel will get a similar outcome as waterfall. [...] [redacted] Do you mean different floors? If yes we need to think the product they [*sic*] should we even allow that.").

bids depend on the floor price and on its probabilistic assessments of other bids. When the advertiser's value is *high* ($2.20), it may find it optimal to bid $1.80 if the floor price is $1 and to bid $2 if the floor price is $2. When its value is *low* ($1.80), it might be optimal to bid $1.50 if the floor price is $1 and not bid at all in an auction with a floor price of $2. If the advertiser multi-homes and bids in two exchanges (Exchange A and Exchange B), then a publisher that solicits bids from those two exchanges for a single impression could exploit the advertiser's bidding behavior by setting a floor price of $2 on Exchange A and a lower floor price on Exchange B. If the advertiser is unable to detect that its bid requests from Exchange A and Exchange B are actually for the same impression, then:

a. When its value is high, it might win the impression on Exchange A with a bid of $2 even when its lower bid of $1.80 on Exchange B would otherwise have won the impression.

b. When its value is low, the advertiser wins the impression on Exchange B.

The effect is similar to that of multi-calling as discussed in Section V: the publisher extracts more revenue from the advertiser than it could with any fixed floor price.

534. Gaming floor prices in this way damages the safety and simplicity of Google's platform because the advertiser cannot rely on the floor price reported at an exchange being the publisher's lowest acceptable price for the impression. Although the example above might suggest that price-fishing is a mere transfer of surplus from the advertiser to the publisher, the problem is worse: total surplus can be reduced due to the changes in advertiser behavior that price-fishing incentivizes. If the advertiser is aware of the

publisher's tactics, its best response is to bid less in an auction with a higher floor price, understanding that the same impression may be available via a different exchange at a lower price. However, such optimization is difficult, particularly because it is in the publisher's interest to hide from the advertisers that they are being called multiple times. One consequence would be the need for advertisers to invest in technology to track impressions and floor prices across exchanges to protect against the possibility of self-competition. If that proved too difficult or costly, advertisers might reduce multi-homing, bid less on *all* impressions, or reduce the number of impressions they bid on.[1061] Thus, by preventing price-fishing (and removing an incentive for advertisers to reduce their bids), UPR also benefited the majority of publishers that did not engage in such tactics.

C. Unified Pricing Rules Protected Advertisers and Benefited Publishers in the UFPA

535. Along with its launch of the Unified First Price Auction in 2019, Google introduced Unified Pricing Rules (UPR).[1062] Before the introduction of UPR, publishers could set floor prices in AdX for specific bidders, including Google Ads, DV360, non-Google DSPs, and ad networks. These floor prices were implemented via "pricing rules," with publishers able to configure up to 5,000 floor prices in this way.[1063] On the other hand,

[1061] *See* Deposition of [redacted] at 236:15-18 [redacted] GOOG-AT-MDL-007177040, at -276 ("You know, if there were [multiple] floors being set to artificially raise prices from an auction standpoint, that could scare buyers away.").

[1062] *See* Jason Bigler, "An update on first price auctions for Google Ad Manager," Google Ad Manager Blog (May 10, 2019), https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/ ("In addition to impacting how publishers are using floor price rules, changing to a first price auction in Ad Manager requires a change in how our rules function. […] That's why we released a new feature to all publishers globally, called unified pricing rules.").

[1063] *See* Lucie Laurendon, "Google Unified First Price Auction Explained," Smart Ad Server Blog (captured on Mar. 7, 2022) at

publishers could *not* use GAM to set floor prices for Open Bidding exchanges and other indirect sources of demand: instead, they needed to configure floor prices for each exchange separately using that exchange's user interface.[1064] The consequence of this piecemeal system was that "[p]ubs set different floors for the same buyer on different exchanges,"[1065] and, as a result, multi-homing advertisers—those who bid across multiple different channels—could face different floor prices for the same impression on different exchanges. As described above, these advertisers "struggle[d]" to de-duplicate impressions and coordinate bids made through different channels and, as a result, could be exposed to price-fishing tactics.[1066]

536. Under UPR, the same floor prices publishers set in GAM apply equally to AdX, other participating exchanges, and remnant line items (including any header bidding demand). This protected advertisers from publishers' price-fishing tactics. UPR did not prevent

---

http://web.archive.org/web/20220307203150/https://smartadserver.com/articles/google-unified-first-price-auction-explained/ ("Publishers will only be able to set UPR with a limit of 200 (vs 5,000 OA rules before).").

[1064] *See* Comms Doc, "Ad Manager Unified 1st Price Auction" (Sep. 27, 2019), GOOG-DOJ-09714662, at -664 ("[Publishers] are unable to [set pricing floors in the Ad Manager UI] for Exchange Bidding and other indirect sources of demand trafficked through non-guaranteed line items."); Declaration of N. Korula (Aug. 4, 2023), GOOG-AT-MDL-008842393, at ¶ 40 ("However, publishers were unable to use the Google Ad Manager user interface to set pricing floors for Open Bidding partners and other indirect sources of demand trafficked through non-guaranteed line items. Instead, publishers had to undertake the complex and time-consuming task of configuring pricing floors separately on each exchange and network where their inventory was available.").

[1065] Presentation, "DRX Unified Yield Management Strategy Review" (Jul. 9, 2018), GOOG-DOJ-11781854, at -869 ("Pubs set different floors for the same buyer on different exchanges to simulate a real-time waterfall and soft floor the buyers (like DBM), and AdX primarily bears the brunt of these higher floors").

[1066] *See* Presentation, "Unified 1st Price Auction" (Mar. 4, 2019), GOOG-DOJ-06525908, at -915 ("Today, buyers struggle to optimize when bidding across different channels due to lack of symmetry: different auction rules and different floor prices can apply for the same impression."). *See also* Sarah Sluis, "Header Bidding Unleashed a Huge Infrastructure Problem and Ad Tech Will Either Sink or Swim," AdExchanger (Apr. 24, 2017), https://www.adexchanger.com/platforms/header-bidding-unleashed-huge-infrastructure-problem-ad-tech-will-either-sink-swim/ ("Though DSPs are seeing a flood of new impressions, many of them are selling the same thing. A publisher with three header bidding partners will have three exchanges sell its inventory, tripling the amount of impressions a DSP must evaluate and tripling the listening cost. […] [M]any DSPs are devoting resources to deduplicating impressions to avoid spending millions in server fees. […] But not all DSPs have the resources to smartly filter out low-value impressions.").

publishers from configuring floor prices based on the advertiser, the ad size, and the inventory type (*e.g.*, display, mobile, and app), with a limit of 200 floor pricing rules at any time.[1067]

537. By protecting advertisers from price-fishing, UPR also benefited publishers. Although individual publishers might be incentivized to engage in price-fishing tactics, advertisers might end up bidding less on *all* impressions, which in turn could end up harming *all* publishers. This is a well-known economic phenomenon akin to the *tragedy of the commons*[1068]: price-fishing publishers impose an *externality* that harms advertisers and other publishers. By preventing publishers from engaging in that type of gamesmanship, UPR protected advertisers and publishers that were not price-fishing.

538. A third-party survey of publishers in February 2020, after the introduction of UFPA with UPR, found that only 4% of respondents described UPR as having a negative impact on their business.[1069] Post-launch experiments found that the changes had [redacted] for non-Google DSPs buying via AdX and [redacted] total publisher revenues [redacted].[1070]

[1067] *See* Google, "Unified pricing rules," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/9298008 ("You can apply up to 200 unified pricing rules per Ad Manager network.").

[1068] *See, e.g.*, Ostrom, E. (2008). Tragedy of the commons. In S. N. Durlauf & L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*, 2nd ed. (pp. 360-363). Palgrave Macmillan.

[1069] *See* Advertiser Perceptions, "SSP Report: Part of the Programmatic Intelligence Report" (Apr. 16, 2020), GOOG-DOJ-AT-00608572, at -573 ("Sample: Digital sales and operations contacts from The Advertiser Perceptions Ad Pros proprietary community and trusted third-party partners as needed [...] Fielded 2/3 to 2/14 2020"), -577 (Figure [dark blue bars correspond to "Negative Impact"]). *See also* Sarah Sluis, "Google Ad Manager Policy Changes Don't Hurt Publishers, According to Advertiser Perceptions," AdExchanger (May 5, 2020), https://www.adexchanger.com/platforms/google-ad-manager-policy-changes-dont-hurt-publishers-according-to-advertiser-perceptions/ ("Google's recently changed rules around unified pricing only negatively impacted 4% of publishers.").

[1070] *See* Presentation, "Changes to Ad Manager, AdMob auction" (Sep. 3, 2019), GOOG-DOJ-14549757, at -772 [redacted] Presentation, "1PA Impact Post Launch" (Nov. 27, 2019), GOOG-DOJ-14566285, at -288 [redacted]

D. Under UPR, Publishers Could Still Preference Demand Sources

539. Even though UPR prevented publishers from setting exchange-discriminatory reserves, publishers still had other tools to favor their preferred demand sources.

540. *First*, a publisher could favor a header bidding exchange by modifying the value CPMs on the line items used to represent that exchange's header bid.[1071] Inflating the value CPMs in that way would make the header bidding exchange more likely to be allocated the impression in the UFPA and would not increase the price for the impression paid by the header bidding exchange (because, as discussed in Paragraph 115 above, the value CPMs reported by publishers do not affect the prices paid by remnant line items). Publishers could also choose to offer some impressions only to their preferred exchanges via header bidding.[1072]

541. *Second*, publishers could also favor certain exchanges or advertisers by offering them post-auction discounts. Post-auction discounts are contractual agreements between publishers and ad buyers and exchanges, in which the publisher agrees to rebate a fraction of the winning bids made by the ad buyer or exchange after the auction has concluded. An ad buyer with such an agreement is incentivized to bid more in each auction for an impression, making it more likely that it wins the impression. As a result, post-auction discounts can be used to increase the share of impressions that are allocated

[redacted]"), -293 [redacted]").

[1071] *See, e.g.*, Asmaâ Bentahar, "Bid Adjustments Simplified: Run Fair Auctions with no Hassle," Pubstack (May 2, 2021), https://www.pubstack.io/topics/bid-adjustments-simplified ("Alternatively, this next one will slightly increase all returned CPMs, giving Prebid an edge in the competition against GAM").

[1072] *See, e.g.*, Prebid.org, "Running Prebid.js without an ad server" (accessed Sep. 7, 2023), https://docs.prebid.org/dev-docs/examples/no-adserver.html ("This example demonstrates running [a header bidding] auction and rendering without an ad server.").

to a given ad buyer or exchange. For example, suppose the publisher agrees to rebate 50% of the winning bids from a certain DSP. If there is a uniform floor price of $3 for an impression and (absent the discount) the DSP would have only been willing to pay $2 for the impression, then *with* the discount, the DSP would be willing to bid $4, which exceeds the floor price. If its bid is the highest bid for the impression, the advertiser bidding using that DSP would be allocated the impression, paying only half its bid to the publisher. As a result, the post-auction discount has a similar effect as a lower floor price for the bidder, but without leading to other harms (such as price-fishing).

542. Post-auction discounts are routinely used in the market for display advertising. An internal Google document from June 2020 notes its use [redacted].[1073] As Michael McNeeley, VP of Product at Index Exchange, described in an interview with Digiday, "[p]ost-auction discounts give publishers' sales team the chance to go and incentivize spend with advertisers just like they do for other forms of media," and, as commented by the article's author, "[t]his is actually not a new trend, as bid prioritization and bid modifiers for certain agencies have been used by publishers in SSP platforms for many years."[1074]

543. In an auction, post-auction discounts are usually a better tool for publishers than differential floor prices. In theory, when bidders are asymmetric in their willingness to pay or in their quality, then an optimal handicapping process causes different bidders to

[1073] "Agency Programmatic Buying Models and Discounts," (Jun. 24, 2020), GOOG-DOJ-AT-00173317, at -318 ("Post-Auction Price Reduction is live with [redacted].").

[1074] Seb Joseph, "Ad-buying agencies are cozying up to SSPs, creating more transparency questions," Digiday (Mar. 31, 2020), https://digiday.com/marketing/ad-buying-agencies-are-cozying-up-to-ssps-creating-more-transparency-questions/.

be charged different prices when they win.[1075] Appropriately designed post-auction discounts can often approximate the optimal process, while exchange-discriminatory floor prices, which use the same prices for any winning bidder whose bid exceeds its floor, cannot generally approximate the optimal mechanism.[1076] As a result, exchange-discriminatory floor prices are a less effective tool.

E. Responding to Plaintiffs' Allegations

1. Plaintiffs Neglect the Industry-Wide Acceptance of UPR and Its Benefits to Advertisers and Publishers

544. Plaintiffs fail to recognize the benefits of UPR to multi-homing advertisers and non-multi-calling publishers in the UFPA. As I have described above, UPR made it easier for publishers to set floor prices for non-Google exchanges in Open Bidding. Maintaining the ability to set exchange-discriminatory floor prices would likely have both made it more difficult for advertisers to bid optimally on Google's platform and led to externalities from price-fishing that would harm other publishers.

545. While providing no concrete evidence, Plaintiffs suggest that publishers would have benefited from setting higher reserve prices "for Google exchange and buying tools."[1077] Yet, they neglect the fact that unified pricing rules are now an industry-wide "best-practice."[1078] For example, Xandr now recommends—as part of its "Seller Best

[1075] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 24-25, 149-53.

[1076] In the technical notes in Section XV.G, I show using an example that the publisher can obtain strictly higher revenues with post-auction discounts than with differential price floors.

[1077] Fourth Amended Complaint ¶ 453.

[1078] *See* Xandr, "Seller Best Practices" (May 5, 2023), https://docs.xandr.com/bundle/industry-reference/page/seller-best-practices.html.

Practices" listed on its "Industry Reference" Documentation Center—to "[e]stablish consistent price floors for the same inventory in all systems (*i.e.*, ad server unified pricing rules, line item CPMs, deal prices, SSP floors, etc.)."[1079]

546. Other online display advertising intermediaries have similar rules requiring floor prices to be uniform across bidders. For example, Meta enforces a "code of conduct" for the auctions it participates in, requiring that all bidders in each auction face the same floor prices,[1080] and [redacted].[1081] As a VP at The Trade Desk expressed, "[i]n a perfect world [...] floor[s] would be consistent and constant" and "would be the same regardless of path."[1082]

547. Professor Gans concludes that, "if Google had not had market power in the publisher ad server market and had not been vertically integrated from that market into an adjacent vertical segment (the exchange market), it would have neither had the ability nor the incentive to engage in the conduct described,"[1083] but the use and preference for such

[1079] Xandr, "Seller Best Practices" (May 5, 2023), https://docs.xandr.com/bundle/industry-reference/page/seller-best-practices.html.

[1080] Meta, "Code of Conduct," Meta Audience Network" (accessed Dec. 2, 2023), https://www.facebook.com/audiencenetwork/partner-program/code-of-conduct ("When a reserve price is applied as part of the auction, it should apply identically to all demand sources."). *See also* [redacted]

[1081] *See* [redacted]

[1082] Will Doherty, "The door in the floors: Transparency, price discovery, and market efficiency in programmatic," The Current (Nov. 8, 2023), https://www.thecurrent.com/will-doherty-transparency-programmatic-data.

[1083] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 458.

rules, including by smaller firms and ones representing just one side of the market, contradicts this conclusion.[1084]

2. Plaintiffs Neglect How UFPA Changed Publishers' Incentives to Set Differential Floor Prices

548. Plaintiffs and their experts argue that "[h]istorically, publishers set different price floors for different exchanges and different buyers in the publisher ad server," and often spent "considerable resources" to do so.[1085] Professor Pathak claims that these higher floor prices were used "to account for the perceived lower ad-quality of impressions served through AdX"—a claim for which he offers no evidence.[1086] And Plaintiffs fail to account for the more obvious reasons that optimal floor prices were higher for AdX before the transition to the UFPA. With earlier ad allocation processes, bids from different exchanges were evaluated sequentially (as in waterfall and Dynamic Allocation) and, often, AdX was running a second-price auction alongside Header Bidding's first-price auctions (as in Open Bidding). In such cases, it was optimal for publishers to set different floor prices for different exchanges, as I explain below. However, those reasons for different floor prices were eliminated by the introduction of the Unified First Price Auction (UFPA).

549. Let us see how setting exchange-discriminatory floor prices could help publishers maximize revenues in a system where the winning bids from multiple exchanges were

---

[1084] *See* [redacted]

[1085] Fourth Amended Complaint ¶ 452.

[1086] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 157.

evaluated sequentially. For example, consider a publisher that sequentially calls two exchanges in the waterfall: Exchange A before Exchange B. Suppose that the publisher believes that the distributions of advertiser values at both exchanges are identical and that the values of bidders at each exchange are statistically independent.[1087] Then the publisher's revenue-maximizing floor price for Exchange A is *higher* than that for Exchange B.[1088] This is because, by choosing a higher floor price for Exchange A, the publisher can try to extract a higher price from the first exchange, knowing that it can offer the impression to Exchange B if it fails to sell on Exchange A.[1089] Similarly, exchange-specific (and impression-specific) floor prices could improve publisher revenues when header bidding exchanges and AdX were not called simultaneously and used different auction rules, as I discussed in Section X.

550. These potential benefits of setting floor prices that differ by exchange are eliminated in the UFPA because bids from all exchanges are evaluated *simultaneously* and *uniformly*. In the above example, the publisher maximizes its revenue in the UFPA by setting the same floor price for both exchanges. This is because the optimal floor price for an exchange trades off the *expected benefits* of increasing the auction's clearing price (when that exchange provides the largest bid and that bid is the only one above the floor price) against the *expected costs* of failing to sell the impression (when that exchange provides

[1087] That is, knowing the value of a bidder at Exchange A does not change the publisher's beliefs about the values of bidders at Exchange B, and vice versa. This is the "independent private values" class of auction models.

[1088] To illustrate, suppose that each exchange has two bidders whose values are drawn independently and uniformly between $0.00 and $1.00. Then the optimal floor price for Exchange A is $0.71 while that for Exchange A is $0.50. In a Unified First Price Auction, the optimal floor price for both exchanges is $0.50.

[1089] *See, e.g.*, Despotakis, S., Ravi, R., & Sayedi, A. (2021). First-price auctions in online display advertising. *Journal of Marketing Research*, *58*(5), 888-907, at "Example 1," 895.

the largest bid and that bid is below the chosen floor price). This tradeoff is the same for the two identical exchanges, and so the optimal floor price is the same as well.[1090]

551. Plaintiffs and their experts fail to account for the fact that publishers would optimally set a higher floor price for AdX prior to the introduction of UFPA even if AdX advertisers had the same quality and distribution of values as other advertisers. In that case, once the transition to UFPA had been completed and advertisers and publishers adapted to the new rules, the earlier reasons to set a higher reserve for AdX would be eliminated and equal reserve prices would be optimal. Even if differences in perceived quality among demand sources are important to some publishers—as claimed by Professor Pathak[1091]—Plaintiffs and their experts do not explain why the filtering tools and the pricing of sensitive categories available within UPR are not sufficient to manage such concerns.

552. In asserting that if "AdX faced the highest reserve pre-UPR, [then] UPR would naturally benefit Google's ad exchange AdX both in win rate and in revenue,"[1092] Professor Weinberg fails to account for the reason that AdX would have faced a higher optimal floor price prior to the transition to the UFPA—the sequential nature of the previous auction format. Even if a publisher rationally set higher floor prices for AdX *before* the UFPA, one cannot assume (as Professor Weinberg does) that it would be optimal for AdX to face a higher floor price *after* the transition to the UFPA.

1090 This finding continues to be true if the number of bidders differ by exchange, since the optimal floor price in the independent private values setting does not depend on the number of bidders (only on the distribution of values associated with those bidders). *See* Riley, J. G., & Samuelson, W. F. (1981). Optimal auctions. *American Economic Review*, *71*(3), 381-392.

1091 Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 157.

1092 Expert Report of M. Weinberg (Jun. 7, 2024), at ¶ 182.

3. Plaintiffs Ignore Alternative Ways Publishers Could Favor Non-Google Demand Sources and Manage Ad Quality After UPR Was Introduced

553. Plaintiffs and their experts claim that UPR prevented publishers from using exchange-discriminatory floor prices for various legitimate business reasons, including "to diversify the sources of demand for their inventory," "to combat [] the problem of adverse selection," and "to improve the quality of ads returned to their site."[1093] Yet, Plaintiffs and their experts fail to mention that publishers had access to more effective tools to achieve these same objectives if the need arose.

554. For example, a publisher wishing to preference a non-Google exchange "to diversify the sources of demand for their inventory" or to "combat problems of adverse selection" could do so by providing a post-auction discount (as discussed in Section XIV.D) or (for header bidding exchanges) by modifying the value CPMs on the line items used to represent that exchange's header bid. Either of these adjustments would lead the favored demand source to win a larger share of the impressions.

555. A publisher wishing to address ad quality concerns also has more direct tools than exchange-discriminatory floor prices. Professor Pathak offers as an "example, a website that publishes stories for children would want to refrain from showing ads related to tobacco or alcohol products."[1094] Firstly, GAM prohibits non-age appropriate ads from being shown to teens and children, and ads for alcohol and tobacco products are even

[1093] Fourth Amended Complaint ¶ 453. *See also* Expert Report of P. Pathak (Jun. 7, 2024), at ¶¶ 157 ("Historically, publishers set higher reserve price floors for AdX to account for the perceived lower ad-quality of impressions served through AdX and increase diversity of demand sources."), 159 ("By eliminating publishers' ability to set differential price floors across exchanges, Google removed a key tool used by publishers to maximize the yield on their inventory and ensure acceptable quality advertisements were displayed on their web pages.").

[1094] Expert Report of P. Pathak (Jun. 7, 2024), at ¶ 171.

more broadly restricted.[1095] In addition, advertisers "can mark [their] ad requests to be treated as child-directed. The feature is designed to help facilitate compliance with the Children's Online Privacy Protection Act (COPPA)."[1096] More broadly, GAM makes ad content protection features available to all publishers, allowing them to avoid presenting ads with sensitive content, ads from specific buyers, or ads in general categories such as "Apparel, Finance, and Health."[1097] Finally, GAM also allows publishers to set category specific floor prices for the sensitive categories.[1098]

556. Publishers could also set advertiser-specific floor prices to address ad quality concerns from specific advertisers. Advertiser-specific floor prices would be more effective than exchange-discriminatory floor prices due to the prevalence of advertiser multi-homing.[1099] When a publisher raises the floor price of a DSP, any low-quality advertisers using that DSP would have an incentive to move their campaigns to another

---

[1095] *See* Google, "Ad-serving Protections for Teens," Google Ad Manager and Ad Exchange Program Policies (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/12171027?hl=en&sjid=2354124647482356745-EU.

[1096] *See* Google, "Tag an ad request for child-directed treatment (TFCD)," Google Ad Manager and Ad Exchange Program Policies (accessed Jun. 24, 2024), https://support.google.com/admanager/answer/3671211?hl=en&ref_topic=28145.

[1097] *See* Google, "Block general categories," Google Ad Manager Help (accessed Jul. 17, 2024), https://support.google.com/admanager/answer/2913554?sjid=2617898862681288879-EU ("You can block high-level groupings of ads — such as Apparel, Finance, and Health — from appearing on your network or specified inventory.").

[1098] *See* Google, "Unified pricing rules," Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/9298008 ("You can set pricing rules that apply only to creatives in selected sensitive categories. Some ads are considered 'sensitive' due to the nature of the business or ad—such as Sensationalism or Significant Skin Exposure. Our system classifies ads automatically, and we don't rely on advertiser-provided categorization.").

[1099] In a 2021 survey, respondent advertisers and ad agencies (who all spent a minimum of $1M annually on digital ads) used an average of 3.4 DSPs and planned to use 5.9 DSPs the following year. *See* Advertiser Perceptions, "DSP Report: Demand-Side Platforms" (2021), GOOG-DOJ-AT-02524665, at -666, -670. *See also, e.g.*, [redacted]

DSP. Only by setting high floor prices for that advertiser *across all exchanges* could a publisher effectively block ads from that advertiser. As discussed in Section XIV.D, publishers could also use post-auction discounts to favor specific exchanges or advertisers with higher-quality ads.

557. Plaintiffs allege that "Google's Unified Pricing rules ensure that rival exchanges and buying tools are at a price disadvantage[] [b]ecause Google's publisher ad server imposes extra fees to serve ad inventory sold on non-Google exchanges [...]."[1100] This is incorrect: no additional fees are charged on ad inventory sold to header bidding exchanges. For exchanges participating in Open Bidding, Google's Open Bidding revenue share pays for valuable services, which for publishers include reporting, payment processing, and integration with non-Google exchanges, and for exchanges include real-time processing of the huge number of bids they submit on each impression. Plaintiffs fail to account for these valuable services when assessing any "disadvantage" to non-Google exchanges. Moreover, a publisher seeking to avoid the Open Bidding revenue share could use header bidding to process real-time bids from relevant exchanges. Or, if a publisher wished to account for fee differences in Open Bidding, it could offer a post-auction discount to Open Bidding exchanges to offset the fee differential.

558. Plaintiffs and their experts also claim that setting exchange-specific price floors could increase the publisher's revenue when the distribution of bidder values differs across the exchanges. Professor Weinberg illustrates his claim with the following example: "imagine that there is a single AdX bidder whose value is distributed uniformly on [12,16], and there is a single OpenX bidder whose value is distributed uniformly on

[1100] Fourth Amended Complaint ¶ 461.

[10,12]. In this case, the revenue-optimal personalized reserves are $13 on AdX and $10 on OpenX."[1101] I show in the technical notes in Section XV.F that the publisher can do *even better* using a uniform reserve price of $12 and a $2 post-auction discount for OpenX. The proper lesson from Professor Weinberg's example is that differences in values can be better managed by post-auction discounts, such as I have described in my published work, rather than by non-uniform reserves.[1102]

559. Finally, Professor Gans claims that "[t]he ability to set flexible pricing floors is valuable to publishers, as it enables them to extract value from high-quality impressions [...]."[1103] He elaborates, "Google buying tools have more information about users than third-party buying tools. As Google's buying tools were more likely to identify high-demand impressions, setting higher pricing floors for Google's buying tools allowed publishers to extract a larger share of the value of those high-demand impressions."[1104] But Professor Gans does not provide any evidence that Google Ads and DV360 actually have "more information" than other buying tools and, even if they do in some circumstances, Professor Gans fails to show the prevalence of such situations and fails to analyze alternative instruments like post-auction discounts that can serve a similar purpose as differential pricing floors.

[1101] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 248.

[1102] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 234-37 (discussing bidding credits as a way to increase revenues with asymmetric bidders).

[1103] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 509.

[1104] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 512.

4. Plaintiffs Fail to Recognize that Each UPR Allows for Up To 50 Advertiser-Specific Floor Prices

560. Plaintiffs' experts claim that advertisers were harmed by the low number of pricing rules, but provide no evidence that was the case. Professor Gans states, "Analyzing Google-produced data on AdX transactions, I find that the average publisher using AdX transacts on average with [redacted] advertisers per month. Large publishers (transacting at least 1 million impressions per month) deal on average with over [redacted] advertisers per month. These numbers greatly contrast with the 200 pricing rules limit imposed by UPR."[1105] This is wrong in several ways. First, as a practical matter, advertisers can be usefully categorized into groups, making it unnecessary to set a separate rule for each advertiser. Second, "pricing rules" are different from "price floors": publishers can "specify up to 50 advertisers per pricing rule," which in principle allows specific floor prices for up to 10,000 advertisers.[1106] Third, Google has policies in place to grant extensions beyond the 200 rules.[1107]

561. Google also demonstrated responsiveness to advertiser concerns about the limits. Based on publisher feedback that 100 rules would be too limiting, Google "increas[ed] the maximum number of pricing rules" to 200 when the UFPA with UPR was fully

[1105] Expert Report of J. Gans (Jun. 7, 2024), at ¶ 510.

[1106] Google, "Unified First-Price Auction - Best practices," Google Ad Manager (accessed Jun. 15, 2024), https://services.google.com/fh/files/misc/unified_first-price_auction_best_practices.pdf; *see* Google, "System maximums and limits," Google Ad Manager Help Center (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/1628457?hl=en.

[1107] *See* Email from [redacted], "Re: UPR 500 rule exceptions" (Nov. 19, 2019), GOOG-DOJ-14028590, at -590 ("Extension requests for < 300 rules -> Lightweight approval [...] Requests for extensions with no rules specified will be capped at 300 [...] Extension requests for > 300 rules -> Needs PM to review business rationale and approve [...] We have generally been granting exceptions whenever a clear reason to justify granular targeting was presented[.]").

launched.[1108] I have not seen any evidence of widespread demand for Google to increase the 200-rule limit further.

[1108] Jason Bigler, "Rolling out first price auctions to Google Ad Manager partners," Google Ad Manager (Sep. 5, 2019), https://blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/.

XV: TECHNICAL NOTES

A. Technical Notes for Section IV (Google Ads Bidding Programs)

1. Theorem 1: Statement and Proof

562. **Theorem 1:** Suppose that Google changes from direct bidding to a bid optimization program, causing its bids to increase and the publisher to revise its floor price.[1109] Further, suppose that the combined effect is that Google Ads' win rate for a given advertiser increases. In addition:

a. Let v denote the advertiser's value for impressions, assumed to be a random variable with finite mean drawn from a twice continuously differentiable distribution F with a strictly positive probability density function f supported on a closed subset of $[\underline{v}, \infty)$.

b. Let λ denote the inverse hazard rate function for the distribution F, defined by $\lambda(v) = \frac{1-F(v)}{f(v)}$.

c. Let M denote the change in the Google Ads win rate.

d. Let S denote the change in the advertiser's *ex ante* expected surplus.

e. Let r denote the publisher's floor price for Google Ads under direct bidding.

If $\lambda(v)$ is non-decreasing for all $[r, \widetilde{v}]$ and $M > 1 - F(\widetilde{v})$, then $S > 0$.

[1109] For this theorem, I assume that other advertisers' bids are unchanged. The profit-maximizing bids for DSPs submitting one bid in the AdX second-price auction would be unaffected by any changes in Google Ads' bids.

563. ***Proof.*** I use the Revelation Principle, conducting my analysis in the space of so-called "direct revelation mechanisms," which map from the advertiser's reported value to interim allocations and payments.[1110] Let $x: [\underline{v}, \infty] \rightarrow [0, 1]$ be the interim allocation function of the advertiser under direct bidding (*i.e.*, $x(v)$ is the probability that an advertiser with value v wins the auction), and let $x': [\underline{v}, \infty] \rightarrow [0, 1]$ be the interim allocation function of the advertiser under the bid optimization program and applying the revised publisher floor price. The increase in the Google Ads win rate after the two changes is then

$$M = \int_{\underline{v}}^{\infty} [x'(v) - x(v)]\, dF(v),$$

which is assumed to be positive for this theorem.

564. By the Milgrom-Segal Envelope Theorem,[1111] for any interim allocation function $y(v) \in [0, 1]$ and interim payment function $t(v)$, the surplus of an advertiser with value v is

$$U(v) = \max_{s \in [\underline{v}, \infty]} \{vy(s) - t(s)\} = \int_{\underline{v}}^{v} y(s) ds.$$

To derive this expression, I used the fact that the bidder with the lowest possible value always loses and pays zero, so $t(\underline{v}) = 0$. Note that

$\int_{\underline{v}}^{\infty} U(v) dF(v) \leq \int_{\underline{v}}^{\infty} (v - \underline{v}) dF(v) < \infty$ because the mean of valuation exists. Then the *ex ante* expected surplus of such an advertiser is

[1110] *See*, *e.g.*, Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58-73.

[1111] Milgrom, P., and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, *70*(2), 583-601.

$$\int_{\underline{v}}^{\infty} U(v)dF(v) = \int_{\underline{v}}^{\infty}\int_{\underline{v}}^{v} y(s)ds\, dF(v) = \int_{\underline{v}}^{\infty}\int_{\underline{v}}^{\infty} 1\{v' < v\}y(v')dv'dF(v) = \int_{\underline{v}}^{\infty} y(v)(1 - F(v))dv.$$

This implies that the change in *ex ante* expected advertiser surplus caused by the bid optimization program is

$$S = \int_{\underline{v}}^{\infty}[x'(v) - x(v)](1 - F(v))dv = \int_{\underline{v}}^{\infty}[x'(v) - x(v)]\lambda(v)\, dF(v)\,.$$

565. Let r' denote the advertiser's *value threshold*, that is, the value that, under the bid optimization program, leads to a bid equal to the publisher's chosen floor price.[1112] There are two cases: either $r' < r$ or $r' \geq r$. That is, the value threshold may be lower or higher than before the changes.

566. In the first case ($r' < r$), since the bid optimization program increases bids, $x'(v) \geq x(v)$ for all v, with strict inequality for some open set of values. Hence the advertiser's *ex ante* expected surplus increases as well:

$$S = \int_{\underline{v}}^{\infty}[x'(v) - x(v)](1 - F(v))\, dv > 0.$$

567. In the second case ($r' \geq r$), the interim allocation functions x and x' satisfy

a. $x'(v) \geq x(v) \geq 0$ for $v \geq r'$, since the bid optimization program increases the bid for an advertiser with value v, and for bids above the relevant floor price, that increases the probability its bid beats those of other bidders.

b. $x'(v) = 0 \leq x(v)$ for $r \leq v < r'$, since all advertisers with values above r had a chance of winning the impression under direct bidding, but only advertisers with

[1112] To make this concrete: under Bernanke with a high bid multiplier βv and post-Bernanke reserve R, the resulting value threshold is $r'=R/\beta$.

values above r' can win the impression after the optimization program is introduced.

c. $x'(v) = 0 = x(v)$ for $v < r \leq r'$, that is, no advertisers with value less than r wins with or without the optimization program.

568. Define the function Δ: $[\underline{v}, \widetilde{v}] \to R$ by

$$\Delta(v) = \int_{v}^{\widetilde{v}} [x'(s) - x(s)] dF(s).$$

The previous paragraph implies that $\frac{d\Delta}{dv}(v) =- [x'(v) - x(v)]f(v)$ is nonnegative for $v < r'$ and nonpositive for $v > r'$, so $\Delta(v)$ is quasiconcave in v (that is, the function is nondecreasing and then nonincreasing on its domain).

569. The condition $M \geq 1 - F(\widetilde{v})$ guarantees that:

$$\Delta(v) = \int_{\underline{v}}^{\widetilde{v}} [x'(v) - x(v)]\, dF(v) = M - \int_{\widetilde{v}}^{\infty} [x'(v) - x(v)]\, dF(v) \geq M - [1 - F(\widetilde{v})] > 0,$$

where the first inequality follows because $[x'(v) - x(v)] \leq 1$. Moreover, because $[x'(v) - x(v)] \leq 0$ for $v < r'$.

$$1 - F(\widetilde{v}) < M \leq \int_{r'}^{\infty} [x'(v) - x(v)]\, dF(v) \leq 1 - F(r').$$

Thus, $\widetilde{v} > r'$ which implies that $[x'(v) - x(v)] \geq 0$ for all $v \geq \widetilde{v}$ with strict inequality for an open set of values. Therefore,

$$S > \int_{\underline{v}}^{\widetilde{v}} [x'(v) - x(v)] \lambda(v)\, dF(v).$$

570. The final step of the proof is to show that $\int_{\underline{v}}^{\tilde{v}} [x'(v) - x(v)]\lambda(v)\, dF(v) \geq 0$. To do so, I use the Barlow-Proschan Lemma[1113], stated below

> **Barlow-Proschan Lemma:** Suppose that W is a signed measure on the interval (a, b) with $\int_{v}^{b} dW(s) \geq 0$ for all $s \in (a, b)$ and that g is a nondecreasing, nonnegative function defined on the same interval. Then $\int_{a}^{b} g(s)dW(s) \geq 0$.

I use this lemma with $dW(s) = (x'(s) - x(s))dF(s)$ and $g(s) = \lambda(s)$ for $s \in(\underline{v}, \tilde{v})$. The condition $\int_{v}^{\tilde{v}} [x'(s) - x(s)]dF(s) \geq 0$ for $v \in(\underline{v}, \tilde{v})$ holds because, as shown above, $\Delta(\underline{v}) \geq 0$, $\Delta(\tilde{v}) = 0$ and $\Delta(v)$ is quasiconcave. The condition that g is a nondecreasing, nonnegative function holds by the assumption of the Theorem.

571. **Corollary 1:** Suppose that the assumptions of Theorem 1 continue to hold, and the distribution of v satisfies DHR for values above r.[1114] Then Google Ads' advertiser surplus also increases: $S > 0$.

[1113] Barlow, R. E., & Proschan, F. (1975). *Statistical theory of reliability and life testing: probability models* (Vol. 1). New York: Holt, Rinehart and Winston.

[1114] That is, that is, the so-called "hazard rate" function $v \mapsto f(v) / [1-F(v)]$ is decreasing in v. Equivalently, DHR is equivalent to the property that the logarithm of the survival function (the so-called "log-survival" function) of distribution F, which is the function $v \mapsto \log[1-F(v)]$, is convex (*i.e.*, the line drawn between any two distinct points on that function's graph lies above the function's graph).

2. Verifying the DHR Condition of Corollary 1 and Calculating Minimum Advertiser Surplus Using Google Ads Data

572. I now use the **Google Ads Log-Level Dataset**[1115] to assess how closely the distributions of advertiser values in Google Ads fit distributions in the family covered in Corollary 1. If the DHR condition holds exactly, Theorem 1 implies that Google Ads' bid optimization programs increased advertiser surplus if the program also increased Google's Ads' win rate. If the DHR condition holds approximately, I calculate a minimal increase in Google Ads' win rate that also guarantees that Google Ads advertisers gain from the programs.

573. Each observation in the Google Ads Log-Level Dataset represents an advertiser's bid in the Google Ads auction. The dataset contains all bids from a random 10% sample of internal auctions conducted between January 17, 2024 and January 23, 2024, inclusive. Altogether, this dataset contains bids for around [redacted] US AdX impressions.[1116] For each impression, the dataset contains estimates of the valuations of all bidders in the Google Ads internal auction, stored in the column "original_unadjusted_score_usd".

574. Because Corollary 1's technical condition applies to the true distribution of advertisers' values—which is not observable using the sample of data I have access to—I compare the *empirical* distributions of advertisers' values in those samples to value distributions that I *know* to be DHR in the relevant domain. To incorporate heterogeneity in

[1115] Google Ads Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000000000 to -000258388.

[1116] This statistic was calculated using code/misc_queries.py in my supporting materials, and the output is saved in code/logs/misc_queries.txt.

advertisers' values for different types of impressions, I examine **slices** of data in the Google Ads Log-Level Dataset, consisting of a group of Google Ads advertiser valuations for a given operating system, platform, browser, domain, GFP network ID, and inventory unit path.[1117] A slice is kept if it contains at least 100,000 bids with advertiser valuations and $1 of advertiser spend across those bids. For each slice, I evaluate the empirical distribution at 100 quantiles, so that each quantile has at least 1,000 affiliated observations.[1118]

575. For each slice, I identify a DHR distribution that closely approximates the empirical distribution of advertisers' values on the relevant set of values (namely, those above the publisher's optimal floor price under direct bidding). To do so, I first note that a distribution has decreasing hazard rate if and only if the log-survival function, $\log(1-F)$, is convex.[1119] Thus, to approximate the empirical distribution, I compute the **convex envelope** of the empirical log-survival function: the pointwise maximum convex function that lower-bounds the empirical log-survival function on the values of interest. I then find the distribution whose log-survival function is equal to this convex envelope.

576. To implement this process computationally, I first find the optimal floor price for the empirical value distribution by computing the revenue a publisher would earn by setting each possible floor price and choosing the one with the highest revenue. I then compute the log-survival function $\log(1-F)$ of the empirical distribution, keeping only the

[1117] Google Ads Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000000000 to -000258388.

[1118] These slices were generated using code/gads_bid_optimization_data.py and saved in code/gads_bid_optimization_data.json in my supporting materials. The number of slices is logged in code/logs/gads_bid_optimization_data.txt.

[1119] This is because the derivative of the log-survival function is the negative of the hazard rate: d/dx log(1-F(x)) = -f(x)/[1 - F(x)].

points between the optimal floor price and the 99th percentile.[1120] I find the convex hull of these points using a standard algorithm, keeping only the lower portion of the convex hull.[1121] Finally, I invert the transformation $F \mapsto \log(1 - F)$ to obtain a fitted DHR distribution.

577. To analyze the goodness-of-fit of the fitted DHR distribution, which I denote by $\hat{F}$, I compare it to its empirical counterpart F. For every value of *v*, $F(v) \leq \hat{F}(v)$ by construction. The area between the two curves over the support of *v* is a measure of distance called the Wasserstein metric or the "earth mover's distance." If the area between the two curves is large, the fitted DHR distribution fits the empirical distribution poorly. Conversely, a small area indicates a good fit.

578. On the vast majority of the data slices, it is possible to identify a DHR distribution that is practically indistinguishable from the empirical distribution of advertisers' values. This means that—for advertisers bidding on those kinds of impressions, which make up the vast majority of Google Ads advertisers—if Google Ads' bid optimization programs increased Google Ads' win rates, they must also have increased the surplus of Google Ads advertisers, even after accounting for the effects on publisher floor prices.

579. To illustrate how well the vast majority of the data slices on Google Ads' values can be fit to DHR distributions, I ordered each data slice according to my measure of its goodness-of-fit normalized by the optimal floor price under direct bidding.

[1120] I exclude the top 1% of values in this analysis because they are sampled too sparsely to accurately assess the shape of the true distribution in this region.

[1121] *See* Scipy, scipy.spatial.ConvexHull, SciPy v1.11.4 Manual (accessed Jan. 10, 2024), https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.ConvexHull.html.

580. In Figure 16, I plot the empirical distribution of advertisers' values in the Google Ads Log-Level Dataset (in blue) and its fitted distribution satisfying DHR (in orange) for nine slices at the 2nd percentile of this goodness-of-fit measure. On these slices, the empirical distribution and the fitted distribution are practically indistinguishable (the blue curve hides the orange curve almost completely), and 98% of data slices fit better than those illustrated.

Figure 16: Empirical Distributions and Their Fitted DHR Distributions for Nine Slices at the 2nd Percentile[1122]

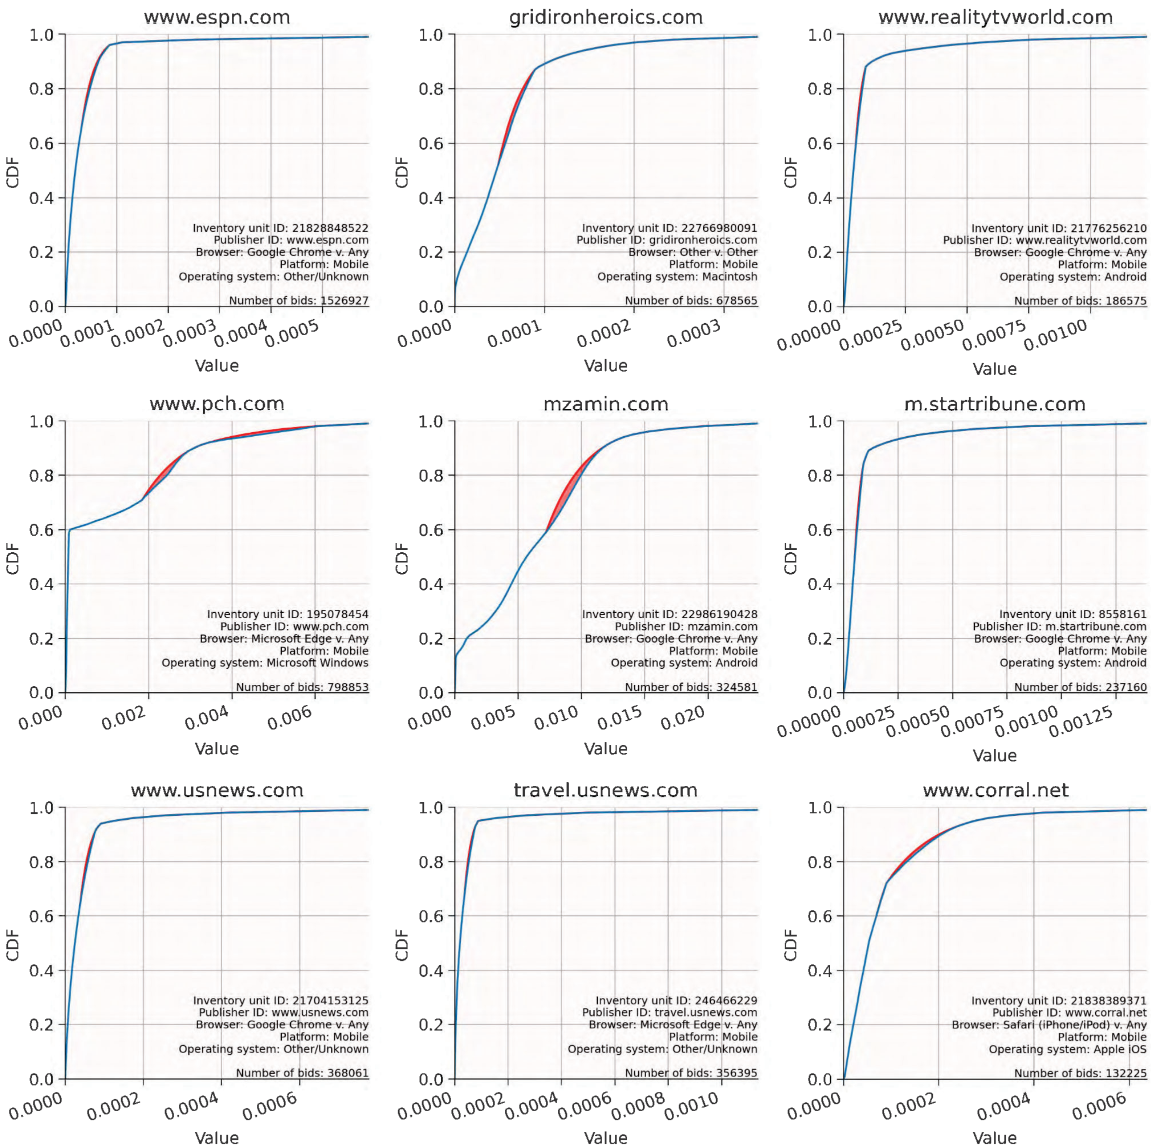

581. I provide similar plots for the nine *worst*-fit slices in Figure 17, showing that while there is some discrepancy between the empirical distribution and its fitted distribution and

[1122] The code to calculate the fits and generate Figures 16 & 17 can be found in code/gads_bid_optimization_fit.py. The figures can be found in code/figures and are prefixed with gads_bid_optimization_grid.

although the orange curve is no longer completely hidden by the blue curve in these plots, the fit is still extremely good.

Figure 17: Empirical Distributions and Their Fitted DHR Distributions for Nine Slices at the 0th Percentile

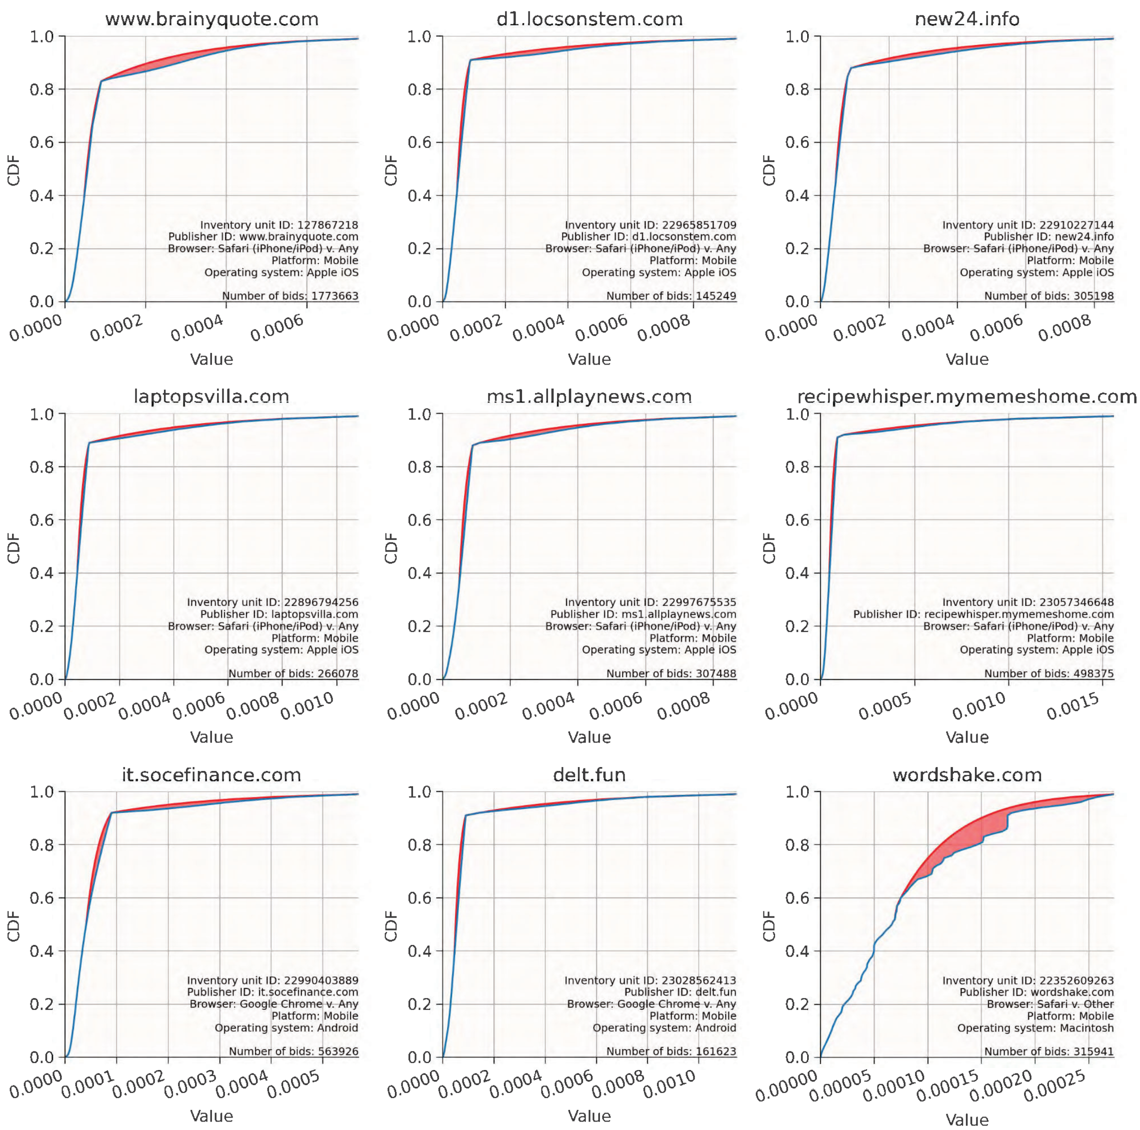

582. For each slice, define M^* as the solution to the following optimization problem:

$$M^* = max_{x(v),x'(v),r'} \int_r^{\bar{v}} [x'(v) - x(v)]\, dF(v),$$

subject to:

a. $S = \int_r^{\bar{v}} [x'(v) - x(v)]\lambda(v)\, dF(v) \le 0,$

b. $x(v)$ and $x'(v)$ are non-decreasing functions,

c. $x(v) \in [0, 1]$, $x'(v) \in [0, 1]$ for all $v \in (r, \bar{v})$,

d. $x'(v) = 0 = x(v)$ for $v < r \le r'$,

e. $x'(v) = 0 \le x(v)$ for $r \le v < r'$,

f. $x'(v) \ge x(v) \ge 0$ for $v \ge r'$.

583. If the first restriction does not bind, it is possible to attain $M^* = 1 - F(r)$, with $r' = r$, and $x(v) = 0$, $x'(v) = 1$ for all $v \in (r,\bar{v})$. However, $S = \int_r^{\bar{v}} \lambda(v)\, dF(v)$ will be strictly positive unless $F(r) = 1$. Thus, if $F(r) < 1$, the first restriction binds and $M^* < 1 - F(r)$; otherwise, $M^* = 0$.

584. M^* is the minimum change in the Google Ads win rate that guarantees a non-negative change in the advertiser's *ex ante* expected surplus, because for all lower values of M it is feasible to achieve a negative surplus.

585. To approximate M^*, I take a grid of K values v_1, v_2, … , v_K, where v_k corresponds to the k/K quantile of the empirical distribution of valuations in the slice and denote by $\tilde{F}(\cdot)$ the

empirical distribution of these K values. Therefore, $F(v_k) = \widetilde{F}(v_k) = k/K$ for all $k = 1, 2, \ldots K$. I define $a = \widetilde{F}(r) \times K$, $b = \widetilde{F}(r') \times K$, and $x(v_k) = \sum_{i=1}^{k} y_i$, $x'(v_k) = \sum_{i=1}^{k} y'_i$ and $R_k = \sum_{i=k}^{K-1} \lambda_i$, where $\lambda_i = (K - i)(v_{i+1} - v_i)$.[1123] I replace $F(v)$ by the empirical distribution of the the K values $v_1, v_2, \ldots, v_K$ in the formulas for M and S, which yields the following optimization problem:

$$M^* \approx max_{b,y,y'} \sum_{i=b}^{K} \frac{K-i+1}{K} y'_i - \sum_{i=a}^{K} \frac{K-i+1}{K} y_i ,$$

subject to:

$$\sum_{i=b}^{K} R_i y'_i - \sum_{i=a}^{K} R_i y_i \leq 0,$$

$$b \geq a,\ 0 \leq y_i \leq 1 \text{ and } 0 \leq y'_i \leq 1 \text{ for all } i,$$

$$\sum_{i=a}^{K} y_i \leq 1,\ \sum_{i=b}^{K} y_i' \leq 1,$$

$$\sum_{i=a}^{h} y_i - \sum_{i=b}^{h} y_i' \text{ for } h \geq b.$$

For a fixed pair (a, b), this is a linear program with inequality constraints.

[1123] Note that this expression converges to the inverse hazard rate as the grid of values becomes dense.

586. To solve the optimization problem for each slice, I set the floor price r to maximize $[1 - F(r)]r$, which is known as the optimal or Myersonian floor price. I set $a = \widetilde{F}(r) \times K,$ and solve the program above solving a linear program for a grid of values $b = a, a + 1,..., K$ and choosing the maximum solution.

587. When we set $K = 100$, in more than 95% of the data slices the increase in win rate that guarantees an increase in advertiser surplus, M^*, is less than 2%. This implies that the advertiser surplus in those slices must increase for an increase of 2% in win rate. The results for other values of M and K displayed in Figure 4.

588. This analysis suggests to me that, after accounting for the incentives for advertisers to adjust their bids and for publishers to adjust their floor prices in response to the programs, the vast majority of advertisers using Google Ads would have benefited from any of Google Ads' bid optimization programs that increased the Google Ads win rate.

B. Technical Notes for Simulations of Dynamic Allocation in Section VIII.E

1. Data

589. I use two main datasets in this simulation analysis.

a. *Google Ads Log-Level Dataset*:[1124] Each observation in this dataset represents an advertiser's bid in the Google Ads auction, which determines the bids that Google Ads will submit to AdX. The dataset contains all bids from a random 10% sample of internal auctions conducted between January 17, 2024 and January 23, 2024, inclusive. After filtering this dataset to observations related to US impressions for

[1124] Google Ads Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000000000 to -000258388.

which a bid was submitted into GAM, this dataset contains bids for around [redacted].[1125] For each impression, the dataset contains estimates of the valuations of all bidders in the Google Ads internal auction.

b. *Google Ad Manager Log-Level Dataset*:[1126] Each observation in this dataset represents a demand source's bid for an impression in the GAM Unified First Price Auction. The dataset contains all bids from demand sources in the same time period as the Google Ads Log-Level Dataset, in those auctions with a viewed winning candidate in the GAM UFPA.[1127] Altogether, this dataset contains bids for around [redacted].[1128] Each observation includes the *inventory unit* of the impression for sale (which is an identifier of the space on a publisher's website where the ad will be displayed), the amount the publisher would be paid if that bid were to win,[1129] and the floor price that the demand source faced when bidding for that impression. By applying the bid inversion technique with modifications as described in the Technical Notes in Section XV.B.4.b to these data, I estimate these bidders' valuations.

[1125] My filtering only includes auctions with bids submitted into GAM (*i.e.*, where publisher_ssp is null and there is an internal auction winner). Of these auctions, 64 of them have multiple internal auction winners; I exclude such auctions from my analysis. Both the number of impressions and this statistic are computed in code/misc_queries.py in the supporting materials and is logged to code/logs/misc_queries.txt.

[1126] Google Ad Manager Log-Level Dataset (January 2024), GOOG-AT-EDTX-DATA-000276098 to -001116097.

[1127] Note the differences in the dataset sampling criteria: 1) the Google Ads dataset is a 10% sample, while the GAM dataset is not, and 2) the GAM dataset contains only auctions with a viewed candidate, while the Google Ads dataset does not have that restriction.

[1128] The number of impressions was calculated using code/misc_queries.py in my supporting materials, and the output is saved in code/logs/misc_queries.txt.

[1129] For remnant line item candidates, this amount may not necessarily be what the publisher is ultimately paid. *See* Letter from J. Elmer to J. Hogan (Mar. 31, 2022), GOOG-AT-MDL-007334120, at -121 ("In addition, for remnant line items, data for publisher payout is entered by publishers and does not necessarily correspond to the actual or agreed publisher payout amount, if the remnant line item wins.").

2. Selecting the Sample of Data from the GAM Log-Level Dataset

590. The GAM Log-Level Dataset contains many observations that are not relevant to my study of DA (for example, because they correspond to pre-auction transactions or correspond to impressions served to users outside of the relevant geographical region for this case). To remove irrelevant bids and auctions from my analysis sample, the GAM dataset is filtered as follows:[1130]

   a. Auctions are limited to those where the end user had a country code of "US" and the field "is_youtube" is False.[1131]

   b. All bids with win_loss values other than "Won" or "Lost" are removed, as those were bids rejected from consideration on the basis of price in their respective auction.

   c. Only line items that directly compete in the GAM unified first-price auction are kept, specifically those with transaction_types of "Open Bidding" and "Open Auction".

   d. Any remaining bid that is below its floor is removed. In the rare case that an auction's winning bid does not exceed its own floor, I discard the auction.

591. For each of the auctions remaining in the dataset, the bid associated with each demand source participating in the auction is calculated as p_view * (pub_payout_usd /

[1130] The data filtering and sample selection steps described in this section are done in code/gam.py in my supporting materials.

[1131] For this and subsequent references to code fields, I use the interpretations of the data fields in the Letter from D. Pearl to K. Garcia (Oct. 6, 2023), GOOG-AT-MDL-C-000012826.

max(views, 1)) and saved in a new column[1132] and bids are rounded to six decimal places, as is consistent with the precision of bids in GAM.[1133] If a demand source places multiple bids into the auction, only one bid is kept: the winning bid if the one of these bids won, and otherwise the highest of the losing bids.

592. For each bidder in each remaining auction, its **price-to-beat**—the minimum amount the bidder needed to bid in order to win the auction—is calculated as the maximum of the bidder's floor and the highest bid by an opponent and saved in a new column.

593. The remaining auctions are grouped by a **triple**, consisting of the publisher domain, inventory unit, and floor price associated with the impression in that auction. In order to ensure that each unit of my analysis is large enough to reliably conduct my simulations, a triple is included in the scope of the study if (i) it contains at least three "eligible" non-Google DSPs who each participated in at least 100,000 auctions and each accounted for at least 0.33% of the revenue in that grouping, and (ii) Google Ads participated in at least 100,000 auctions and accounted for at least 1% of the revenue in that grouping.

594. Heterogeneity is notable in these data: different publishers and inventory units experience different patterns of demand. In addition, because the auction run by GAM is a first-price auction, the bid for an impression might depend on floor prices, both for strategic reasons and because the floor price may reflect additional factors known to the publisher and bidders. To accommodate heterogeneity without additional assumptions, I obtain separate

[1132] This calculation is consistent with the [redacted] (Oct. 6, 2023), GOOG-AT-MDL-C-000012826, at -875, which notes that [redacted]

[1133] Letter from [redacted] (Dec. 7, 2023), GOOG-AT-MDL-C-000012885, at -893 ("[redacted]

empirical estimates for each triple. I analyzed [redacted] inventory units across [redacted] publishers, comprising approximately [redacted] US publisher real-time bidding revenue in the January 2024 GAM data sample.[1134, 1135] I conducted a separate simulation for each of the resulting [redacted] triples (publisher, inventory unit, floor price).

3. Selecting the Sample from the Google Ads Log-Level Dataset

595. I collected pairs of highest and second-highest Google Ads values from all Google Ads internal auctions in which the query ID is contained in the GAM Log-Level Dataset and Google Ads bids into GAM (so that I could subsequently match observations between the datasets).

596. I required two values from Google Ads in my analysis: those associated with the winning internal bid and the internal bid with the highest non-winning unadjusted score.

597. For each bid, the advertiser's value is calculated by dividing the column "original_unadjusted_score_usd" by 1000. If there is no losing unadjusted score, the value for the highest losing bidder is taken to be zero. I remove all pairs where the highest value exceeds the 99.9th percentile value, because the calculated values above that quantile appear to be anomalies in the data.[1136]

[1134] In this statistic and others related to total publisher revenue, I exclude impressions won by "Reservation" and "HBYG," as I understand these aliases to potentially represent numerous, distinct sources of demand. Reservation creatives are associated with remnant line items, while HBYG is an alias for any header bids integrated using GAM's "Header Bidding Yield Groups" feature. *See* Letter from [redacted] GOOG-AT-MDL-007334131, at -134 ("Reservation refers to remnant line items."); *see also* George Levitte, "Improved header bidding support in Google Ad Manager," Google Ad Manager (Apr. 27, 2022), https://blog.google/products/admanager/improved-header-bidding-support-in-google-ad-manager/.

[1135] These statistics are computed in [redacted]

[1136] For example, some observations in the data have values over [redacted] CPM, which seem to be anomalies in the data set of advertisers' values. After applying the filtering described above, no values over $[redacted] CPM remain. This

4. Modeling Bidder Values

598. My model assumes that conditional on the observed triple of characteristics of the impression—the publisher, the inventory unit, and the floor price—bidders' values are statistically independent.[1137] This is a weaker assumption than the "independent private values" assumption made by Professor Weinberg[1138]: any model with independent private values must satisfy this conditional independence assumption. This assumption is consistent with standard practice for modeling display advertising auctions in the academic literature.[1139]

a) Values for AdX

599. For my simulations, I model bidders on AdX as consisting only of the bidders on Google Ads. In reality, many bids on AdX come from DV360 and Authorized Buyers, but I make this simplifying assumption for two reasons. *First*, I can *directly* observe a distribution of values of Google Ads advertisers, because I have access to these data for a subset of

restriction has the effect of reducing the revenue accruing to publishers from AdX, leading to a conservative estimate of the publisher's revenue from DA.

[1137] While I would not expect bidders' values for an impression to actually depend on the impression's floor price, allowing bidders to condition on the floor price captures some of the characteristics of the impression that are observable by bidders *and* publishers, but not included in the data.

[1138] Expert Report of M. Weinberg (Jun. 7, 2024), at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...] and it is a sensible assumption to make").

[1139] *See* Balseiro, S. R., & Gur, Y. (2019). Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, *65*(9), 3952-68, at 3956 ("The information provided by the auctioneer heterogeneously affects the value advertisers perceive for the impression on the basis of their targeting criteria. The values advertisers assign to the impression they bid for […] are assumed to be independently distributed across impressions and advertisers"); Choi, H., & Mela, C. F. (2023). Optimizing reserve prices in display advertising auctions. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4523022*,* at 10 ("Following the symmetric independent private value assumption commonly adopted in prior work (*e.g.*, Ostrovsky and Schwarz 2011, Balseiro et al. 2015), advertiser valuations are assumed to be drawn independently and identically from the conditional distribution $F_V(v|z)$, where Z are auction specific observed covariates.").

impressions in the sample.[1140] This means that for Google Ads bidders, I have accurate estimates of values without needing to perform the "bid inversion" step that I describe below. *Second*, the Google Ads dataset also allows me to see the *two* highest values among bidders from Google Ads data. The second-highest bid on AdX is needed in order to simulate the second-price auction used with DA on AdX.[1141] Because this assumption omits some bidders, it underestimates the revenue accruing to publishers from AdX in each simulation and hence underestimates the revenues received by publishers under DA. I provide additional details about my approach to simulating AdX values in the Technical Notes in Section XV.B.3.

600. For each simulation, I first "flip a coin" that is weighted to match the participation probability of Google Ads observed in the log-level data for the corresponding triple of publisher-inventory unit-floor price.[1142] Conditional on the outcome of that coin toss, I draw a random pair of Google Ads valuations from the Google Ads Log-Level Dataset.

b) Bid Inversion to Estimate Values for Non-Google Demand Sources

601. For non-Google demand sources, I observe bids instead of values in the GAM log-level auction data, so I perform bid inversion to estimate the value distribution of the three largest eligible non-Google demand sources (by revenue) for each inventory unit. I allow the estimated value distributions to vary between demand sources.

1140 [redacted] *See* Letter from [redacted], (Jun. 9, 2023), GOOG-AT-MDL-C-000012751 ("[redacted]").

1141 I do not need data on the second-highest bids on other demand sources, because I do not model the internal auctions that may or may not be run on those demand sources, only the *net* bids they make.

1142 I calculate the percentage of auctions in the triple for which Google Ads places a bid. Then, for each auction that I simulate, Google Ads draws a value with this probability.

602. I assume that demand sources generate most bids in the GAM data by first observing an advertiser's value and then selecting a bid that maximizes the advertiser's surplus, given bids expected for impressions sharing those characteristics. I assume that each demand source bids with a probabilistic assessment of their price-to-beat for the auction, which is the least bid that the bidder needs to make to win the auction. This probabilistic assessment is determined using a **cumulative distribution function** (CDF) $F(\cdot)$ that tells the demand source for each bid b the probability $F(b)$ that the bid wins the auction, with that CDF estimated using data from previous auctions. I assume that F is differentiable except at the floor price (at which there is an atom), so $F(b)$ is the probability that the maximum of the floor price and the highest other bid is less than or equal to b. The demand source selects a bid b for an advertiser of value v that maximizes the expression $(v - b)\,F(b)$.

603. Together, these assumptions imply that when a bidder bids b its value v for that impression is equal to $b + F(b)/f(b)$, where $f(b)$ corresponds to the probability density function (PDF) of the demand source's historical price-to-beat distribution.[1143] The distribution of bids observed in the GAM auction data can be used to infer $F(b)$ and $f(b)$, making it possible to recover a value that rationalizes each bid. The resulting estimates of the distribution of values do not rely on assumptions about the shape of the distribution of values and are specific to each demand source.

604. This methodology is standard for the empirical analysis of first-price auctions. However, it is also well known in the literature that this method cannot be applied blindly. To make

[1143] Since I assume that the price-to-beat has a differentiable CDF almost everywhere (with an atom at the floor price), the probability density function (PDF) can be calculated using the derivative of the cumulative density function.

sure that the results are robust, I need to apply some standard corrections to the inferred distribution of values in order to obtain bid functions that may be rationalized given the data and economic theory. In particular:

a. *Smoothing*: It is easier to estimate from the data the CDF of bids, $F(b)$, than the PDF, $f(b)$, which is the former function's derivative. In order to estimate this derivative reliably (*i.e.*, without creating spurious data artifacts), it is necessary to "smooth" the observed $F(b)$. I use a standard method for this smoothing called kernel density estimation (KDE).[1144, 1145]

b. *Truncation*: It is difficult to rationalize the very highest bids observed in the GAM data as surplus-maximizing bids by bidders with reasonable values for an impression.[1146] My interpretation of these very highest bids (the top 2% of bids) is that they are mostly associated with experimentation by bidders (or possibly a result of some bidding mistakes or responses to information I do not observe in my sample),[1147] and so, for the highest 2% of bids, I set the bidder value to be

[1144] *See* Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first‐price auctions. *Econometrica*, *68*(3), 525-574; Perrigne, I., & Vuong, Q. (2019). Econometrics of auctions and nonlinear pricing. *Annual Review of Economics*, *11*, 27-54.

[1145] Concretely, KDE spreads the probability mass of each data point in the empirical distribution of bids (log-transformed) over a small neighboring region according to a kernel function that integrates to one. I use the Gaussian density as a kernel function to obtain a smooth estimate of the density of the natural logarithm of bids. Then, by changing variables, I obtain the density of bids. The parameters of the Gaussian distribution are chosen using the inbuilt KDE procedure for SciPy (a popular Python library), which uses a heuristic due to Scott (1992) to choose the so-called *bandwidth* (which affects the standard deviation of the Gaussian kernel). *See* SciPy, "scipy.stats.gaussian_kde," SciPy v1.11.4 Manual (accessed Dec. 3, 2023), https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html; Scott, D.W. (1992). *Multivariate density estimation: Theory, practice, and visualization*. John Wiley & Sons.

[1146] This is because in first-price auctions, very high bids do not substantially increase the probability of winning an auction but do substantially increase the amount that winning bidders must pay. The only way to rationalize very high bids is then with implausibly high values for impressions or by assuming that these bids are formed differently, for example as a result of experimentation.

[1147] [redacted]

equal to the value inverted for the bid at the 98th percentile of the bid distribution. The effect of this value-capping procedure on revenues and fill rates in my waterfall simulations is the same as restricting publishers to set their floor prices so that each demand source fills an impression at least 2% of the time it is called to bid.

c. *Enforcing monotonicity*: Economic theory tells us that surplus-maximizing bids must be *increasing* functions of these bidders' values, yet the estimation procedure described above sometimes results in small regions of values on which the estimated bidding function decreases. To deal with this, I use a procedure to impose monotonicity of the bidding function while still ensuring that each bidder's inferred value results in an observed bid that maximizes its expected profits from bidding.[1148]

605. Upon calculating these values, sampling from a demand source's value distribution is equivalent to first flipping a coin weighted as the participation probability of the demand source in the log-level data, and, conditional on the outcome of that coin, drawing a random value from the estimated values for that demand source.

2016), GOOG-DOJ-AT-02471119, at -120 ("[redacted].").

[1148] Concretely, for each value v on a grid, I find all bids that satisfy the *first-order* conditions for optimality, $v=b+F(b)/f(b)$, and then among those bids, I find the bid $b(v)$ that is the global maximizer of the expected utility $(v-b)F(b)$. Because $(v-b)F(b)$ exhibits increasing differences in values and bids (v,b) for any *CDF* $F(b)$, this procedure ensures that $b(v)$ is monotonic according to Topkis's theorem. *See* Milgrom, P., & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157-180. I then use piecewise linear interpolation to obtain $b(v)$ for values between points in the grid. Only 3.7% of inferred values are changed by more than 10% as a result of this procedure to enforce monotonicity (logged in code/logs/parse_da_results.txt).

5. Simulating Sales of Impressions

606. After obtaining distributions of bidder values for each demand source, I then simulate the sales of impressions using both DA and the waterfall. To simulate such a sale, I randomly draw values for each of the non-Google bidders for the inventory unit from the estimated value distributions. I limit my attention to three non-Google demand sources for each auction, both for computational tractability and because the latency associated with calling more than three or four demand sources sequentially makes longer waterfalls impractical.[1149] For AdX, I draw a *pair* of values from the Google Ads value distribution. A pair of values is needed in order for me to simulate the outcome of the second-price auction used under DA. I adjust those advertisers' values for the Google Ads revenue share to determine bids into AdX. For each triple, I simulate 5,000,000 auctions.

a) Simulating the Waterfall

607. In the waterfall, the publisher calls a sequence of demand sources. In my simulations, when a demand source is called, it is offered the opportunity to buy the impression at a posted price. These posted prices are chosen separately for each demand source to maximize the publisher's expected revenues. If the value drawn for the demand source in the simulation exceeds the posted price, the demand source purchases the impression and pays the publisher the posted price. Otherwise, the request is passed to the next demand source in the waterfall, with the process repeating until the impression is sold or the list is exhausted (leaving the impression unsold). I select an order of the demand sources and

[1149] *See, e.g.*, Bidgear, "A Well Organized Passback Strategy," Bidgear Blog (Aug. 7, 2016), https://bidgear.com/blog/a-well-organized-passback-strategy-24 ("You don't want to make your waterfall chain too big because your site latency will suffer and some ads might time out.").

posted prices for each to maximize the publisher's total expected revenues using a dynamic programming procedure described in the Technical Notes in Section XV.B.6. By assuming that publishers set optimal posted prices, my simulation tends to overstate the revenues of the waterfall.

b) Simulating DA

608. To simulate DA, I assume that AdX first runs a second-price auction with a floor price *R* set by the publisher. If the highest AdX bidder's value for the impression exceeded the floor, the publisher would receive the larger of *R* and the second-highest AdX value; otherwise, the impression would be offered to non-Google demand sources using a waterfall. As in the other simulations, I calculate the optimal orderings and posted prices for the waterfall using a dynamic programming procedure described in Section XV.B.6. Calculating the optimal floor price *R* for AdX in this way is computationally difficult, so I instead assume (conservatively, as explained below) that publishers experiment to choose the floor price of AdX in DA, selecting the *R* from the following set of heuristics that leads to the largest average revenue:

a. The floor price AdX had faced in the counterfactual waterfall (if applicable),

b. The expected revenue from the waterfall,

c. 110% of the expected revenue from the waterfall, and

d. 125% of the expected revenue from the waterfall.

The waterfall described in heuristics b, c, and d refers to the procedure used when the highest AdX bid fails to exceed the floor. I include heuristics c and d because, according

to economic theory, the optimal floor price must be larger than the expected revenue from the waterfall. It is unlikely that any of these floor prices is the publisher's revenue-maximizing floor price for AdX, and switching to an optimal floor price would necessarily lead to higher publisher revenues. My reliance on the four heuristics just described tends to underestimate the benefits of DA for publishers.

c) Comparing DA and the Waterfall: Two Different Counterfactuals

609. I measure the effects of DA under two different counterfactuals:

a. *Counterfactual 1:* In the first counterfactual, I compare calling AdX with DA to a baseline in which the waterfall would *only* call three non-Google demand sources. In this counterfactual, enabling DA brings AdX as a new source of demand for publishers. I compare outcomes under this baseline waterfall to those that would have arisen if AdX were called with DA and, if AdX did not win the impression, the impression would be allocated through a waterfall containing just the two best non-Google demand sources.[1150]

b. *Counterfactual 2*: In the second counterfactual, I compare calling AdX with DA to a baseline in which the waterfall includes both AdX (not using DA) and non-Google demand sources. In this baseline waterfall, AdX and three non-Google demand sources all participate using publisher-optimal posted prices. I compare this baseline to the outcome of calling AdX with DA, followed by a

[1150] I reduce the number of non-Google demand sources in the DA simulation to ensure that any increase in publisher revenues caused by DA is not driven by an increase in the total number of demand sources that the publisher accesses. The counterfactual in which the publisher does not displace an existing demand source is covered in Section VIII.C: with an appropriate floor, DA is a *risk-free revenue improvement* for the publisher.

waterfall with the same three non-Google demand sources. This comparison isolates the effect of DA's introduction of auction-based pricing. This counterfactual is relevant for assessing whether DA improved outcomes compared to a baseline in which all demand sources competed on the same basis, using the waterfall.

6. Calculating Prices and Orders in the Waterfall

610. I calculate the order of each publisher's waterfall list and the prices for each demand source in the waterfall using a combination of search and dynamic programming.

611. Given a fixed ordering, to determine the optimal floor prices, I use dynamic programming. The fundamental challenge is that the optimal floor price for a demand source in the waterfall depends on the floor prices chosen for subsequent demand sources. So, to overcome this challenge, I first determine floor prices for the last demand source, *Z*, and then the second-last demand source, *Y*, and so on. For the last demand source, *Z*, I use a "grid search" to determine the optimal floor price, testing each of a grid of possible floor prices and finding one that maximizes the expected revenue (which equals the product of that floor price and the probability that a demand source has a value for the impression that exceeds that price). The expected revenue becomes the continuation value of not selling the impression to demand source *Y*. I then calculate the optimal floor price for *Y* in the waterfall by conducting another grid search, this time maximizing the publisher's expected revenue from selling the impression to *Y* or *Z*, by

incorporating the continuation value.[1151] Similarly, this maximized expected revenue becomes the continuation value of not selling the impression to the demand source before *Y*: demand source *X*. I continue to calculate floor prices in this fashion until I obtain a floor price for each demand source.

612. I consider each possible ordering of demand sources in the waterfall, calculating the optimal prices associated with that ordering (as described above) and the expected revenue for the publisher. For each waterfall, I choose the ordering of demand sources that results in the highest expected revenues for the publisher.

7. Simulation Results

613. I measured auction outcomes according to two metrics. The first is the average *publisher remnant revenue*. The second is the *match rate*, which is the fraction of impressions that are successfully sold before the waterfall is exhausted. Higher match rates mean that fewer impressions go unsold.

614. I first report the effects of DA under Counterfactual 1.[1152] I found that:

a. [redacted] of simulated publishers experienced revenue increases from DA, with total publisher remnant revenue increasing by [redacted]. The heterogeneity of those effects is plotted in Figure 9 and discussed in Paragraph 290. For the majority of publishers, the gains are substantial, with the median publisher experiencing a [redacted] increase in remnant revenue.

[1151] Mathematically, the objective of that grid search is the proposed floor price times the probability that the demand source's value exceeds that floor price *plus* the probability that the demand source's value does *not* exceed the floor price times the continuation value of not selling the impression.

[1152] All simulation results are computed in code/parse_da_results.py in my supporting materials and logged to code/logs/parse_da_results.txt.

b. DA lifted the overall remnant match rate from [redacted] leaving many fewer impressions unmatched.

615. In Counterfactual 2, I find that:

a. [redacted] of publishers experience revenue increases from DA, with overall publisher remnant revenue increasing by [redacted]. Heterogeneity in these effects is plotted in Figure 10 and discussed in Paragraph 291. The median publisher experienced a remnant revenue increase of [redacted].

b. Across all publishers, DA increased overall remnant publisher revenue match rates, from [redacted]

8. Robustness Checks

616. I ran additional simulations to check the robustness of these results to several of my modeling decisions and data inclusion criteria.

617. To check the robustness of my results to the value capping procedure described in Paragraph 604, I reran my simulations capping values at the 95th percentile and the 99th percentile. Under Counterfactual 1, the resulting increases in revenues under those alternative assumptions are [redacted]. Match rates increased from [redacted] (compared to from [redacted] Under Counterfactual 2, the resulting increases in revenues are [redacted] respectively (compared to [redacted] with the cap at the 98th percentile). Match rates increased from [redacted] nd from [redacted]

[redacted] These results all demonstrate the robustness of the conclusion that DA increases publisher revenues and match rates.[1153]

618. I set my inclusion criteria to include only triples for which there is enough data to reliably estimate values for each demand source. While my criteria allow me to incorporate a large portion of the data, less stringent criteria would have allowed me to achieve even more coverage (at the expense of less reliable estimates). To make sure that my results were not sensitive to my exact definition, I reran my experimental pipeline, this time requiring Google Ads and third-party demand sources to participate in only [redacted] auctions in a triple (instead of [redacted]. Doing so creates more eligible triples: with this less stringent criterion, there are [redacted] qualifying triples, comprising [redacted] of US publisher real-time bidding revenue in the GAM log-level data ([redacted][1154] Using these new triples, under Counterfactual 1, revenue increases by [redacted] and match rates increase from [redacted] Under Counterfactual 2, revenue increases [redacted] and match rates increase f[redacted]).[1155] Again, these results demonstrate the robustness of the conclusion that DA increases publisher revenues and match rates.

619. Another way I could have created the simulated auctions would have been to insert a step after running inversions and before sampling values in which all auctions that do not

[1153] These results were computed using code/parse_da_results.py in my supporting materials, with the numerical values logged in code/logs/parse_da_results.txt.

[1154] These statistics were computed using code/misc_queries.py in my supporting materials, with the numerical values logged in code/logs/misc_queries.txt.

[1155] These results were computed using code/parse_da_results.py in my supporting materials, with the numerical values logged in code/logs/parse_da_results.txt.

contain at least one of the inverted demand sources or Google Ads are removed. Doing so would increase the simulated participation rates of all demand sources, including Google Ads. With this added preprocessing step, under Counterfactual 1, revenue increases by [redacted] and match rates increase from [redacted] Under Counterfactual 2, revenue increases by [redacted] and match rates increase from [redacted][1156] Yet again, these results demonstrate the robustness of the conclusion that DA increases publisher revenues and match rates.

C. Technical Notes for Section IX (Enhanced Dynamic Allocation)

1. Further Details on Calculating the EDA Price to Maximize Publishers' Revenues While Fulfilling Guaranteed Contracts

620. Three major components are key to making EDA work: forecasting, computing the EDA price, and resolving ties between remnant line items. In this section, I focus on the original method EDA used to resolve such ties, called *randomized assignment.*[1157]

621. I illustrate randomized assignment using an example. Consider a publisher with 3,000 similar impressions to be allocated, of which 2,000 must be reserved for guaranteed contracts. When a publisher receives a bid of, say, $1, then to maximize its revenues, it

[1156] These results were computed using code/parse_da_results.py in my supporting materials, with the numerical values logged in code/logs/parse_da_results.txt.

[1157] *See* Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 to -181. As header bidding became more common, publishers increasingly used value CPMs to represent demand sources bidding in real-time. In this environment, EDA also increased expected publisher revenue compared to the pre-EDA procedure. Google later transitioned to a tie-breaking procedure I call "randomized bid perturbation," which could only increase publisher revenues further. *See* Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -182 [redacted]

should accept the bid unless it forecasts that there will be higher bids on at least 1,000 impressions in the future. Google's EDA program helps publishers to achieve this allocation by first forecasting future remnant bids and then using that forecast to estimate the highest price p such that there are at least 1,000 impressions receiving bids of at least p. I will call that p the "opportunity cost of not serving the guaranteed contract" or just the *market-clearing price*.

622. In auctions with multiple identical or similar items for sale, the total number demanded may exceed supply even at the market-clearing price. For example, suppose there are 950 impressions with bids higher than \$0.50 and 75 with bids equal to \$0.50. At any price higher than \$0.50, there is less demand than the supply of 1,000 impressions for remnant bidders, while at a price equal to \$0.50, there are 1,025 impressions demanded, which is 25 more than the available supply. This example highlights a common challenge for auctions with multiple similar units: the need for *tie-breaking*.

623. EDA resolved these ties with a procedure called *randomized assignment (RA)*. To illustrate randomized assignment, suppose that there is a group of 3,000 similar impressions and 2,000 of these need to be allocated to guaranteed contracts. Based on historical data, Google forecasts that [redacted] impressions will receive AdX bids larger than [redacted], and [redacted] of those will be larger than [redacted] Since there are only [redacted] remnant impressions to fill and Google knows that, just from AdX, it will receive at least [redacted] bids larger than [redacted], the market-clearing price is no less than [redacted]. This [redacted] minimum, which is computed using AdX bids only, was known as the EDA price.[1158] I denote it by p_{EDA}.

[1158] *See* Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 ("A distribution of AdX bids is built for each inventory slice in offline

624. Suppose that the highest remnant line item has a static bid of \$1.50. As the bid is static (it is a value CPM), the same bid price is applied to each impression. This means that the market-clearing price of each impression is $max(p_{EDA}, \$1.50)$.

625. EDA accepts all 500 bids from AdX that are larger than the market-clearing price (\$1.50) and must break ties to accept 500 of the 2,500 bids exactly at the price of \$1.50 (from the static remnant line item). [redacted]

626. The general idea is that EDA accepts the AdX bid when it exceeds *both* the remnant bid and p_{EDA}. When the remnant bid exceeds the AdX bid and p_{EDA}, then it wins with *just the right probability so that guaranteed contracts are fulfilled and revenue is maximized (subject to filling all the impressions with some ad).*[1160]

pipeline. Based on p[redacted]").

[1159] *See* Design Doc, "Uniform Treatment for DFP Remnant and AdX under EDA" (Apr. 2019), GOOG-AT-MDL-011687180, at -180 to -181.

[1160] *For interested readers, this note illustrates the computation of "just the right probability so that guaranteed contracts are fulfilled."* Suppose that a remnant line item has a static bid of \$1.50 for all of a publisher's impressions in some group, that a fraction q of those must be reserved for guaranteed contracts, and that a fraction f of AdX bids are forecasted to be less than \$1.50. If $f < q$, then there are more than enough AdX bids higher than \$1.50 to serve all the non-guaranteed impressions, so no impressions are assigned to remnant demand. In that case, the market-clearing price is more than \$1.50. If instead $f > q$, then the publisher can earn more revenue by selling some impressions to the remnant line item: the market-clearing price for the publisher's relevant impressions is \$1.50. AdX bids greater than \$1.50 will win a fraction $1 - f$ of all the impressions. Of the remaining fraction f, q will be reserved for guaranteed contracts. So, to accept only the highest bids, $f - q$ should be assigned to the remnant line item with its bid of \$1.50 and the remaining such bidders should be rationed. To achieve that, when the remnant line item bids of \$1.50 are highest, they must be assigned an impression with probability $(f - q)/f$.

2. *Theorem 2: Statement and Proof*

627. **Theorem 2:** Suppose that publishers' guaranteed contracts are unchanged after the introduction of EDA and that Google accurately forecasts the distribution of future bids from AdX. Then (1) EDA increases the publisher's expected revenue relative to the pre-EDA allocation procedure, and (2) if publishers set the optimal floor price for the AdX auction ignoring direct contracts, the floor set by EDA *maximizes* publisher revenue.

628. ***Proof.*** We first prove part (1) of the Theorem.

629. Let q denote the estimated proportion of impressions that must be allocated to guaranteed contracts and F denote the cumulative distribution function of the highest bid from AdX (that is, $F(p)$ is the probability that the highest bid from AdX is less than or equal to p). I assume that F is continuous and has bounded density.

630. I use v to denote the value CPM of the most competitive non-guaranteed line item and μ for the average payment of this line item. Recall from the main text that $\mu \leq v$, as would be optimal for a publisher. Additionally, let $REV_{AdX}(a)$ denote the expected revenue from AdX *conditional* on the highest AdX bid exceeding its floor price a.

631. The following relevant properties of EDA were true at launch:

   a. The EDA price, p_{EDA}, is set so that $q = F(p_{EDA})$.

   b. The EDA algorithm was implemented as follows:

i. If AdX's bid exceeded the larger of p_{EDA} and v, then AdX won the impression.[1161]

ii. Otherwise, if $v \geq p_{EDA}$, then the impression was allocated to the non-guaranteed line item with probability $\frac{F(v)-q}{F(v)}$ and to the guaranteed contract otherwise.

iii. In any other case, the impression was allocated to the guaranteed contract.

632. The expected revenue per impression from indirect demand sources without EDA is $R_{pre-EDA} = (1 - q)[F(v) \cdot \mu + (1 - F(v)) \cdot REV_{AdX}(v)]$. In this expression, $(1 - q)$ is the probability that an impression is available to non-guaranteed buyers. It is multiplied by the term $[F(v) \cdot \mu + (1 - F(v)) \cdot REV_{AdX}(v)]$, where $F(v)$ is the probability that the highest AdX bid is less than the value CPM. Hence, $(1 - q)F(v)\mu$ is the expected revenue from remnant demand and $(1 - q)(1 - F(v))REV_{AdX}(v)$ is the expected revenue from AdX.

633. I now compute the expected revenue with EDA. There are two cases to consider.

a. First, consider the case where $v < p_{EDA}$. Then, all remnant impressions are sold to AdX, leading to expected revenues of $R = (1 - F(p_{EDA}))\, REV_{AdX}(p_{EDA})$. Since $q = F(p_{EDA})$, R is equivalent to $(1 - q)\, REV_{AdX}(p_{EDA})$. Also,

[1161] To simplify exposition, I assume there is no publisher-set AdX-specific floor price (or that any such floor price is smaller than the EDA price or v). The argument is similar if a larger AdX-specific floor price is present and F is appropriately adjusted to account for this floor.

$REV_{AdX}(p_{EDA})$ is greater than both v (and μ) and $REV_{AdX}(v)$ (since $p_{EDA} > v$), so $R \geq R_{pre-EDA}$ as desired.

b. Second, consider the opposite case where $v \geq p_{EDA}$. In this case AdX wins each impression with probability $(1 - F(v))$ and generates average revenue $REV_{AdX}(v)$ when it wins. A non-guaranteed line item wins with probability $F(v)\frac{F(v)-q}{F(v)} = F(v) - q$ and generates average revenue μ when it wins. In total, this means the expected revenue in this case is $R = (1 - F(v))REV_{AdX}(v) + (F(v) - q)\mu$, which may be rewritten as $R = ((1 - F(v))REV_{AdX}(v) + F(v)\mu) - q\mu$. Now I expand $R_{pre-EDA}$ to obtain $[(1 - F(v))REV_{AdX}(v) + F(v)\mu] - q[F(v)\mu + (1 - F(v))REV_{AdX}(v)]$.

Comparing the two expressions, one can observe that $R \geq R_{pre-EDA}$ if and only if $q(F(v)\mu + (1 - F(v))REV_{AdX}(v)) \geq q\mu$. The latter inequality holds because $REV_{AdX}(v) \geq v \geq \mu$ and hence $F(v)\mu + (1 - F(v))REV_{AdX}(v) \geq \mu$.

Hence, in both cases, the expected revenue from EDA, R, is no less than, $R_{pre-EDA}$.

634. We now show part (2) of the Theorem: if publishers set the optimal floor price for the AdX auction ignoring direct contracts, the floor set by EDA *maximizes* publisher revenue. Since—by assumption—Google accurately forecasts the distribution of bids from AdX, the EDA price is such that the $1 - q$ fraction of impressions allocated to

AdX remnant demand is *exactly* the set of impressions with the highest bids on AdX. This means there is no way to reallocate impressions between remnant demand and direct contracts to increase the bids received on remnant inventory. Additionally, as long as the floor price exceeds the static optimal reserve for that distribution (which is enacted by EDA under the assumption that publishers set the optimal static floor price, since EDA respects that floor as well), there is no way to increase the floor price to increase revenues on those impressions. Since there is no reallocation of impressions that increases bids received for remnant impressions and no change in the floor price that increases the remnant revenues conditional on those bids, there is no way to increase publisher revenues. This means that EDA maximizes publishers' remnant revenues under the assumptions of Theorem 2. ■

3. Analysis of pCTR Data: Supplemental Figures

635. Figure 18 and Figure 19 are analogous to Figure 12 and pertain to pAVR and pVTR, respectively.[1162]

Figure 18: Scatterplot of pAVR by publisher with/without EDA

[1162] These figures were generated by running the file code/pctr.py in my supporting materials. The figures are located at code/figures/pAVR_regression.png and code/figures/pVTR_regression.png, respectively.

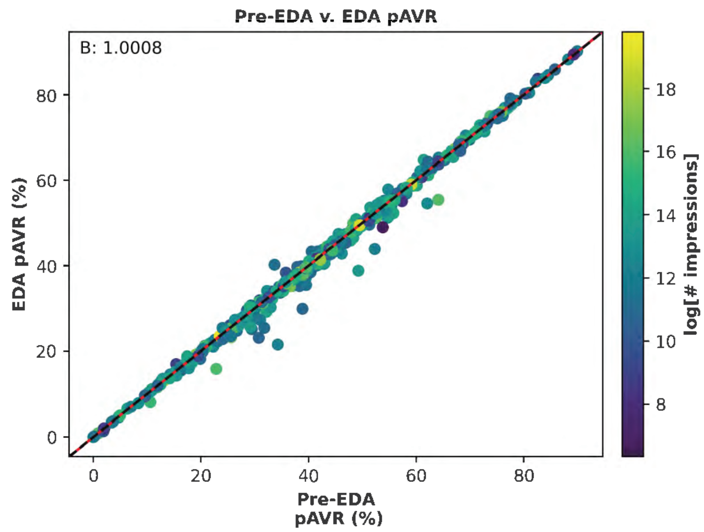

Figure 19: Scatterplot of pVTR by publisher with/without EDA

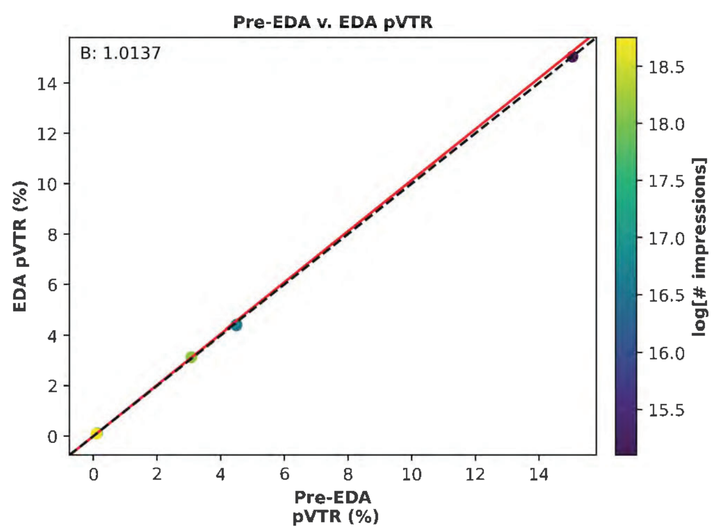

D. Technical Notes for Section X (Header Bidding and "Last Look")

1. Theorem 3: Statement and Proof

636. **Theorem 3:** Suppose that (i) a publisher sells an impression to a fixed set of bidders, including bidders on AdX, (ii) the publisher does not know each bidder's value for the

impression, (iii) bidders do not know each other's values for the impression, and (iv) all participants understand the rules and have the following information:

a. Each bidder knows its own value for the impression.

b. Each bidder's value is drawn from a commonly-known probability distribution[1163] and is statistically independent from other bidders' values.[1164]

c. Each bidder determines a bid as a function of its value to maximize its surplus from the impression, given its probabilistic assessments about the bids of other bidders.

Then, if the publisher chooses revenue-maximizing floor prices, it earns exactly the *same* expected revenue with a single unified second-price auction with all bidders participating as it would with first-price header bidding auction followed by a second-price auction for AdX.[1165] Moreover, given those optimal floor prices, each header bidder has the same chance of winning and the same expected surplus as an identical bidder on AdX.

637. ***Proof.*** This theorem is proved as a corollary of a well-known result from auction theory for this model: the *Payoff and Revenue Equivalence Theorem,* a version of which

[1163] I assume that this probability distribution F with density f satisfies a technical condition called "Myersonian regularity," which requires that the function $v \mapsto v - (1-F(v))/f(v)$ is increasing, and ensures that the optimal mechanism for selling the impression is an auction.

[1164] This implies that each bidder and the publisher can make a probabilistic assessment about other bidders' values and estimates of other bidders' values would not be changed upon learning one bidder's value.

[1165] More specifically, the theorem considers an auction game in which the publisher moves first, the header bidders second, and AdX bidders last. The publisher selects header bidding floor prices and a function to map each header bid it receives into a floor price for the AdX auction; the header bidders select their bids; and then finally the AdX auction is run.

states:[1166] *Suppose that bidders' values are independent and let r be any floor price.*[1167] *If bidders bid optimally, all auctions that award an item to the highest value bidder with bid above r lead to the same expected revenue for the seller and the same expected surplus for each bidder.*

638. By the Payoff and Revenue Equivalence Theorem, which applies under the assumptions of Theorem 3, if the publisher configures header bidding and its floor prices to ensure that the highest value bidder with value above the floor price r wins the auction, then its average revenue from that arrangement will be the same as for a unified second-price auction with floor price r and including all the same bidders.

639. To ensure that the highest value bidder with value above the floor price r wins the auction, the publisher could set a floor price of r in the header bidding auction and commit to choosing floor prices in the AdX auction as a function of the header bids received, as follows. Let $\beta(v)$ be the implied optimal bidding rule for header bidders, which determines for a bidder of value v its best bid $\beta(v)$ in the header bidding auction. It is a standard result that function β is increasing and, by the logic of bid shading into a first-price auction, $\beta(v) < v$.

[1166] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 73-77.

[1167] Roughly, bidders' values are independent if knowing the value of one bidder does not help predict the value of the other bidder. The same result holds if values are independent conditional on commonly observed information about the item. This "independent private value[s]" assumption is maintained by Plaintiffs' experts. *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024) at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...] it is a sensible assumption to make since (a) internal Google documents demonstrate that Google assumes this as well (e.g., GOOG-AT-MDL-004016180), (b) bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's) and since the bidders do not know each other's identities, even if they learn about others' bids it could possibly not be that useful towards determining their own valuation.").

640. There are two possibilities: first, if the highest bid in the header bidding auction is $B > r$, the publisher could estimate that the winning bidder has value $V = \beta^{-1}(B) > B > r$. If the publisher sets a floor price of V in the AdX auction, an AdX bidder wins only if it has the highest value (which is also higher than r). The second possibility is that there is no header bid exceeding the floor price r. In that case, if the publisher sets the floor price to be r for the AdX auction, a bidder wins only if it has the highest value and that value is larger than r.

641. With this floor pricing policy, in both cases, the winning bidder has the highest value among all bidders and that value is above the optimal floor price r. By the Payoff and Revenue Equivalence Theorem, this implies that the publisher obtains the same expected revenue and each bidder obtains the same expected surplus as for a unified second-price auction with floor price r (which is optimal among the class of all mechanisms for selling the impression). ◾

E. Technical Notes for Section XII (Sell-side Dynamic Revenue Sharing)

642. In the following proofs, I assume that the standard sell-side revenue share of 0.8 is applied, otherwise each argument applies identically, replacing 0.8 with the negotiated revenue share.

1. Theorem 4: Statement and Proof

643. **Theorem 4:** If publishers did not change their floor prices and bidders did not change their bids, DRS v1 could only increase the number of impressions sold, publisher revenues, and buyer surplus.

644. ***Proof.*** For each impression, there are three possible scenarios under DRS v1: (1) AdX does not win the impression, (2) AdX wins the impression with its standard revenue share, and (3) AdX wins the impression with a lower revenue share. In the first two scenarios, the winner and their payment are the same with or without DRS v1: buyer surplus and publisher revenues are unchanged. In the third scenario, the impression sells on AdX when it would not have done in the absence of DRS, and the publisher receives revenue equal to the impression's floor price. As long as this floor price is no smaller than any header bidding or remnant demand offer for the impression, this implies that the publisher's revenue weakly increases. In this case, the buyer pays its bid, which can only increase its surplus, because a surplus-maximizing bidder chooses a bid less than its value. If a buyer bids its values—the optimal strategy for a single impression sold via a second-price auction—it would pay its value for the impressions sold by DRS v1, which would lead to a zero effect on buyer surplus. ■

2. Theorem 5: Statement, Proof, and Further Discussion

645. **Theorem 5:** Suppose that a publisher is selling an impression to a fixed set of bidders on AdX. The publisher does not know each bidder's value for the impression, and bidders do not know each other's values for the impression, but all participants have the following information:

a. Each bidder knows its own value for the impression.

b. Each bidder's value is drawn from a commonly known probability distribution and is statistically independent from other bidders' values.[1168]

[1168] This implies that each bidder and the publisher can make a probabilistic assessment about other bidders' values and estimates of other bidders' values would not be changed upon learning one bidder's value.

c. Each bidder determines a bid as a function of its value to maximize its surplus from the impression, given its probabilistic assessments about the bids of other bidders.

Then, if the publisher chooses revenue-maximizing floor prices, it earns a higher expected revenue on an impression to which DRS v1 is applied than it would without the program, and bidder surplus is unchanged.

646. ***Proof.*** This theorem is proved as a corollary of a well-known result from auction theory for this model: the *Payoff and Revenue Equivalence Theorem,* a version of which states:[1169] *Suppose that bidders' values are independent and let r be any floor price.*[1170] *If bidders bid optimally, all auctions that award an item to the highest value bidder with bid above r lead to the same expected revenue for the seller and the same expected surplus for each bidder.*

647. The Payoff and Revenue Equivalence Theorem implies that the exact method for determining payments (*e.g.*, first-price auction versus second-price auction) does not matter to determine the expected revenue—all that matters is the probability each bidder with a given value wins the impression.

[1169] *See* Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press, at 73-77.

[1170] Roughly, bidders' values are independent if knowing the value of one bidder does not help predict the value of the other bidder. The same result holds if values are independent conditional on commonly observed information about the item. This "independent private value[s]" assumption is maintained by Plaintiffs' experts. *See, e.g.*, Expert Report of M. Weinberg (Jun. 7, 2024) at ft. 62 ("I assume for the majority of this report that the advertisers have independent private values for impressions [...] it is a sensible assumption to make since (a) internal Google documents demonstrate that Google assumes this as well (e.g., GOOG-AT-MDL-004016180), (b) bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's) and since the bidders do not know each other's identities, even if they learn about others' bids it could possibly not be that useful towards determining their own valuation.").

648. Let r^* be the optimal floor price on a publisher's impression before the introduction of DRS. Before DRS, the publisher always earns [redacted] of the total revenue from sales, so, if the publisher optimizes the floor price, it must be the case that selling to the highest bidder with value greater than r^*/[redacted] maximizes the *total* revenue from the sales.

649. If the publisher sets a floor price equal to r[redacted] on an impression for which DRS v1 is active, the impression still sells to the highest bidder with value greater than the r[redacted]. The Payoff and Revenue Equivalence Theorem implies that the expected revenue is unchanged, and that this is the optimal floor price for the impression. Bidder surplus is also unchanged.

650. Since the total revenue is the same and AdX takes a lower revenue share on some impressions, the total earnings for the publisher *must increase*. ■

651. The statement and proof of Theorem 5 covers the case in which DRS v1 is active on an impression. In practice, AdX used probabilistic throttling to avoid reducing its overall revenue share too much, and as a consequence, buyers and publishers would need to consider the possibility of throttling when determining their optimal bids and floor prices. One simple solution for buyers would be to reduce their bids into the dynamic region only so often as to avoid being throttled by AdX. Buyers could monitor when AdX started to throttle their bids into the dynamic region and return to their pre-DRS strategies for some time. Another alternative would be for buyers to shade bids less to reduce their downward influence on revenue shares. Adapting bidding strategies in these ways would improve buyers' outcomes, and similar approaches could be taken by publishers. In either case, I would expect buyers and publishers to learn to adapt their strategies over time to

adjust for the possibility of throttling, and that publisher revenues would increase overall as a result of DRS v1, corresponding to Google's experimental results. In practice, AdX used [redacted] throttling probabilities for most buyers.[1171] Throttling may not have been necessary very often if buyers adapted their strategies to the presence of throttling or if many transactions were still cleared with the standard [redacted] revenue share because at least two bids were above the floor price.

3. Theorem 6: Statement and Proof

652. **Theorem 6:** Revenue for AdX under DRS v1 increases *only if* publisher revenues from AdX also increase. Additionally, every percentage increase in revenue for AdX results in a proportionally larger increase in revenue for publishers.

653. ***Proof.*** Let rev be the total revenue from sales before the introduction of DRS and $\widehat{rev}$ be the total revenue from sales after DRS v1. Let $prev =$ [redacted] rev be publisher revenues from AdX before DRS is introduced and $\widehat{prev} =$ [redacted] $\widehat{rev}$ be publisher revenues from AdX after DRS v1. Let $AdXRev =$ [redacted] rev be the AdX revenue before DRS and $\widehat{AdXRev} = 0.19\,\widehat{rev}$ be AdX revenue after DRS v1. Then $\frac{\widehat{prev}}{prev} =$ [redacted] This implies that for every [redacted] increase (or decrease) in AdX profits, publishers experience at least a [redacted]% increase (respectively, decrease) in revenues from AdX. (As noted above, I focus in this section on the [redacted] revenue share, but for publishers with different

[1171] *See* Presentation, "Dynamic Sell-Side Revshare[:] GDN/DRX Summit 2015" (Nov. 2, 2015), GOOG-DOJ-13202659, at -671.

negotiated revenue shares, the numbers appearing in this calculation would be different, although the general result remains unchanged.) ■

4. Technical Description of DRS v2

654. Let v_1 and v_2 respectively denote the highest and second-highest gross bids into AdX, and let bidder 1 and bidder 2 be the respective advertisers submitting these bids. Let r denote the floor price applying to the impression.

655. The auction allocation and pricing worked as follows under DRS v2:[1172]

a. Case 1: $r < v_1 \leq r/$[redacted]. AdX cannot win the auction with a fixed [redacted] per-impression revenue share, so it reduces its revenue share on the impression.

i. Bidder 1 pays $(v_1 + r)/2$ to AdX.

ii. AdX pays r to the publisher.

iii. AdX updates the debt accounts for the publisher and the buyer. Bidder 1's debt increases by $\Delta_b = r/$[redacted]$- (v_1 + r)/2$, the amount it would need to raise its bid to win with a fixed [redacted]% revenue share. The publisher's debt increases by $\Delta_p = r -$ [redacted]$v_1 + r)/2$, the amount the floor price would need to lower by to sell with a fixed [redacted] revenue share.

b. Case 2: $r/$[redacted] $< v_1$. AdX may collect its standard revenue share on the impression and additionally recoup debts previously accrued by publishers and advertisers, while charging Bidder 1 less than its bid and paying the publisher more than the floor.

[1172] *See* Design Doc, "Dynamic Revenue Sharing (DRS) V2 Proposal" (Mar. 24, 2015), GOOG-DOJ-13221355, at -356 to -358; Launch Doc, "AdX Dynamic Revshare v2: Launch Doc" (Jan. 13, 2016), GOOG-DOJ-13207875, at -879 to -881.

i. Bidder 1 pays the standard price $max(v_2, r$/[redacted] plus a debt payment Δ_b to AdX. Δ_b was chosen to ensure that it was never more than the advertiser's outstanding debt balance and the buyer never paid more than its bid v_1 (in fact, never more than half of the buyer surplus in the absence of debt reclamation). Moreover, if the buyer set its *own* price (by submitting a second bid or nonzero "minimum payment"), any payment the buyer made in excess of r/[redacted] would be deducted from its debt. This served as an incentive for buyers to submit second bids.

ii. AdX pays the publisher the standard price max[redacted]$v_2, r)$ plus a share of the advertiser's debt payment [redacted]$_b$ minus a publisher debt payment Δ_p. Δ_p was never more than the publisher's debt balance and was chosen to ensure that the payment to the publisher was at least the floor r.

c. Case 3: if $v_1 \leq r$, AdX does not win the auction and no payments are made.

5. Lemma 1: Statement and Proof

656. **Lemma 1:** Suppose that AdX recoups all debts under DRS v2. Then, buyers accrue the full debt on each impression (equal to the difference between the floor price that would apply in the absence of DRS and the price it pays under DRS v2) and, after accounting for the payment of buyer debt to publishers, publishers accrue zero debt on net in expectation.

657. ***Proof.*** If all debts are repaid, the total debts paid by publishers equal the total debts they accrue on impressions sold via DRS v2 minus [redacted] times the total debts paid (and accrued)

by buyers. But on each individual impression, the debt accrued by a publisher is equal to [redacted] times the debt accruing to a buyer. Therefore, the net debt paid by publishers on impressions sold due to DRS v2 equals zero. Since buyers are never paid part of the publishers' debts, if all debts are repaid, the buyer pays the full debt it accrues on each impression sold due to DRS v2. ■

6. Theorem 7: Statement and Proof

658. **Theorem 7:** If publishers do not change their floor prices and buyers do not change their bids, then DRS v2 can only increase the total number of impressions sold and total publisher revenues compared to the absence of DRS.[1173]

659. ***Proof.*** Prior to DRS v2, publishers sell exactly the impressions where the highest bid v_1 clears the pre-revenue share floor $r/$[redacted] and the publisher earns the larger of r and [redacted]v_2 on each such impression. The revenue earned on each such impression is unchanged after the introduction of DRS v2.

660. With DRS v2, the publisher also earns a revenue of r from AdX on impressions for which the highest bid v_1 is between r and $r/$[redacted] (as argued above, the publisher can ignore the debts accrued which cancel out on average). The effect of these additional sales on the publisher's revenue with DRS v2 depends on the nature of the floor r. If r is exactly the price at which the publisher would otherwise sell the impression to an alternative exchange, its revenue is unchanged. Otherwise, if r is a publisher-set floor price (so that the impression would be otherwise unsold) or if r reflects the optimal floor

[1173] I assume that AdX performs enough transactions that all debts in DRS v2 are resolved.

price determined by the publisher given the header bids received, the publisher's revenue strictly improves with DRS v2. ■

7. Theorem 8: Statement and Proof

661. **Theorem 8**: If buyers and publishers set bids and floors to maximize their payoffs after the introduction of DRS v2, then buyer surplus and publisher revenues are the same as in the absence of DRS.

662. ***Proof.*** As Lemma 1 showed, buyers pay debt on each impression sold by DRS v2 equal to the difference between the pre-revenue share floor ($r/$[redacted] and the price paid for the impression. This means that the effective price of an impression cleared by DRS v2 for an advertiser is the pre-revenue share floor $r/$[redacted] As a consequence, buyers should only seek to win an impression if they value the impression above $r/$[redacted] Because each buyer always pays the larger of the second-highest bid and the floor if its bid is larger than $r/$[redacted] the optimal bidding strategy to bid truthfully.[1174] Consequently, for any fixed floor price, the auction's outcome is the same in the second-price auction as it is with DRS v2, and thus the optimal floor prices under DRS v2 are the same as in the absence of DRS v2. Therefore, once buyers and publishers fully adapt their strategies to DRS v2, the total publisher revenues and AdX profits should equal those quantities *before* the introduction of DRS. ■

[1174] To see this, suppose that there was any positive probability that a bidder bid into the dynamic region. Then any such bidder could increase its probability of winning without changing the price it paid by increasing its bid to r/[redacted]. But for bids of r[redacted] or larger, the auction is the same as a second-price auction, which means that each bidder's optimal strategy is to bid truthfully.

8. Technical Description of tDRS

663. The tDRS program worked as follows:

a. AdX observed the publisher-set floor price and the value CPMs associated with any remnant demand line items. The maximum of these values was AdX's floor price r.

b. Before offering the impression to buyers, AdX compared the pre-revenue share floor price, $r/$[redacted] to the Reserve Price Optimization (RPO) program's reserve r^*, which was the floor price that AdX predicted would maximize the publisher's expected revenues from AdX.

c. If the pre-revenue share floor r[redacted] exceeded the RPO reserve r^* for that impression and AdX predicted that it was unlikely for there to be a buyer with value larger than $r/$[redacted] then AdX sent a floor of r with a 0% revenue share on the impression to its buyers, increasing the probability that a sale will occur.

   i. If there was a bid above the floor r and the second-highest bid was less than $r/$[redacted] then the publisher was paid r and accrued a "debt" of 0[redacted] representing the revenue share foregone by AdX on that impression.

   ii. If there were two bids above $r/$[redacted], then the winning buyer would be charged the second-highest bid and AdX would take its standard [redacted] revenue share: in this case, payments were the same as in the absence of DRS. No "debt" was accrued.

d. Otherwise, if $r/$[redacted] was larger than the RPO reserve r^* but AdX predicted it to be likely that there was a buyer with a larger value than $r/$[redacted] then AdX maintained its [redacted] per-impression revenue share and the impression transacted per usual.

e. Finally, if $r/$[redacted] was less than the RPO reserve r^*, then r^* became AdX's floor price. If the impression sold with this RPO reserve and the publisher had previously accrued debt from a DRS transaction, some debt was repaid on this impression, no larger than the difference between r^* and r.

9. Theorem 9: Statement, Proof, and Further Discussion

664. **Theorem 9:** If publishers adjust their floor prices on AdX to maximize profits after the introduction of tDRS and tDRS accurately predicts buyers' bids, total publisher revenues from all demand sources will be higher with tDRS than with a fixed revenue share.

665. ***Proof.*** Suppose that in the absence of tDRS, a publisher sets a floor r for AdX. Without DRS, the impression is sold if there is some buyer willing to pay at least the pre-revenue share floor, $r/$[redacted]

666. Now suppose that tDRS is introduced. I show that the publisher can choose a floor price that guarantees them at least the same revenue as in the absence of DRS. This implies that if the publisher sets its floor prices optimally, its revenue must be even larger.

667. Suppose that the publisher reports a floor of r[redacted] to AdX, which is the pre-revenue share floor in the absence of tDRS. Then there are two possibilities:

a. tDRS applies a [redacted] per-impression revenue share, which means it predicts that there will be a buyer who can beat the new pre-revenue share floor, $r/$[redacted]. If the tDRS prediction model is accurate, the publisher earns at least $r/$[redacted] for these impressions, more than it did in the absence of DRS.

b. tDRS predicts that there will not be an AdX buyer who can beat the price $r/$[redacted] and applies a 0% per-impression revenue share. In this case, the AdX buyer wins the impression if it bids above the floor price of $r/$[redacted] exactly the same as the pre-revenue share for advertisers before DRS. Because tDRS incentivizes truthful bidding, the AdX buyer wins the impression if and only if it would win it without tDRS. In this case, the payment to the publisher for the impression is r[redacted] and the publisher accrues a "debt" of 0.[redacted] , leading to a net payoff of r. The expected publisher revenue is therefore unchanged by tDRS.

668. Thus, the impression sells with the same probability under tDRS as without DRS, and the publisher earns more revenue than without DRS.[1175] Under the optimal floor price, which may differ from $r/$[redacted] the publisher can only receive even higher revenues. ◾

669. Although an assumption of Theorem 9 is that AdX is perfectly able to predict whether its buyers will clear a given floor price, the same conclusion holds as long as (1) AdX predicts bids better than publishers and (2) the revenue share under tDRS maximizes Google's expected profits (or, equivalently, the expected revenues of the publishers).

[1175] In this analysis, I have fixed the bids of header bidders after the introduction of tDRS. In general, header bidders need to account for the change of format in the AdX auction because it may affect their expected payoff of any given bid. However, because under the posited strategy of reporting floor r/0.8, there is no change in the probability of sale on AdX, the header bidders' bids should not change either. Under the publisher's optimal strategy, however, header bidding advertisers' bids would change, but this does not affect the conclusion of Theorem 9.

F. Technical Notes for Section XIII (Open Bidding)

1. When Publishers Boost Bids in Professor Weinberg's Own Example, Publisher Revenues are Identical to a Unified Auction and Neither Bidder Has an "Advantage"

670. I consider Professor Weinberg's example in Appendix D of his report in which there are two bidders "in the independent private values model" with "values are each drawn independently and uniformly from [0,10]."[1176] In Professor Weinberg's example, the "two bidders participate in a first-price auction, but one bidder [(Bidder One)] learns the others' bids [(Bidder Two)] before bidding."[1177] I show that when the publisher boosts the bids of Bidder Two by a factor of two, it achieves the same revenues as it would in a unified auction, while leaving neither bidder with an "advantage."

671. After Bidder Two's bid is doubled, that boosted bid serves as a floor price in a first-price auction with Bidder One. Hence, in this example, Bidder One maximizes its surplus by submitting a bid equal to the boosted bid of Bidder Two when its value exceeds the boosted bid. Bidder Two determines how much to bid by optimizing its expected surplus, $(v - b)Pr\{v_{Bidder\ One} \leq 2b\}$. Maximizing this expression with respect to b implies that it submits bids equal to $b = v/2$.

672. *First,* notice that Bidder Two wins if and only if its value is larger than Bidder One's value. When combined with the fact that both bidders' valuations are drawn from identical uniform distributions, this implies that each bidder has win probability equal to $\frac{1}{2}$.

[1176] *See* Expert Report of M. Weinberg (June 7, 2024) at Appendix D, ¶ 6.

[1177] *See* Expert Report of M. Weinberg (June 7, 2024) at Appendix D, ¶ 5.

673. *Second,* the publisher achieves revenues of

$$E[v_{Bidder\ 2} \mid v_{Bidder\ 1} \geq v_{Bidder\ 2}]\, Pr\{v_{Bidder\ 1} \geq v_{Bidder\ 2}\} \ +$$

$$E[1/2v_{Bidder\ 2} \mid v_{Bidder\ 2} > v_{Bidder\ 1}]\, Pr\{v_{Bidder\ 2} > v_{Bidder\ 1}\},$$

as Bidder 1 pays the bid of its competitor when it wins, while Bidder 2 pays its own bid when it wins. Simplifying this expression further yields

$$\frac{1}{2}E[v_{Bidder\ 2} \mid v_{Bidder\ 1} \geq v_{Bidder\ 2}] + \frac{1}{2}E[1/2v_{Bidder\ 2} \mid v_{Bidder\ 2} > v_{Bidder\ 1}] = 3\frac{1}{3}\,.$$

674. *Third,* notice that each conditional expectation in the previous expression is equivalent. Namely, that,

$$E[v_{Bidder\ 2} \mid v_{Bidder\ 1} \geq v_{Bidder\ 2}] = 3\frac{1}{3} = E[1/2v_{Bidder\ 2} \mid v_{Bidder\ 2} > v_{Bidder\ 1}].$$

Hence, the bidders' average payments are the same.

675. As Professor Weinberg notes in Lemma 3 of his Appendix D, this revenue is identical to the expected revenue in the unified auction with two bidders.[1178]

2. Extending Professor Weinberg's Example Shows that the So-Called "Last Look" Does Not Inherently Decrease Publisher Revenues, Even When Publishers Do Not Boost Header Bids

676. Consider a simple extension of Professor Weinberg's example, with the following two buyers:

[1178] Expert Report of M. Weinberg (June 7, 2024) at Appendix D, ¶ 13 ("With two bidders whose values are each drawn independently and uniformly from [0,10], the sealed bid first-price auction achieves an expected revenue of 10/3").

a. *Buyer One* is an exchange that commits to running a second-price auction and receives bids from two bidders with independent and uniform values over [0, 10]. This assumption differs from Professor Weinberg's as it adds a second bidder to better understand the revenue effects in a second-price auction.

b. *Buyer Two* is a buyer that has independent and uniform values over [0, 10]. This assumption is identical to Professor Weinberg's.

For the sake of comparison, in this example, I adopt Professor Weinberg's assumption that the publisher does not set a floor price.

677. *First,* consider the case in which Buyer One has no "last look" and passes its internal clearing price to a first-price auction between Buyer One and Buyer Two. Given that Buyer One runs a second-price auction, the bid function for each bidder in Buyer One is $b_{Buyer\ 1,\ Advertiser\ i}(v) = v$. On the other hand, when Buyer Two has a valuation of *v*, it best-responds by submitting the bid *b* that maximizes its expected surplus:

$$U_{Buyer\ 2}(v,\ b) = (v - b)\,(1 - Pr\{Buyer\ 1\ Wins\})$$

$$= (v - b)\,(1 - Pr\{v_{Buyer\ 1,\ Advertiser\ 1} > b,\ v_{Buyer\ 1,\ Advertiser\ 2} > b\})$$

$$= (v - b)(1 - (1 - F(b))^2)$$

$$= (v - b)(1 - (1 - (b/10))^2)\ .$$

For any *v*, the optimal bid for that *v* must satisfy the first-order condition

$\frac{\partial U_{Buyer\ 2}(v,b)}{\partial b} = 0$, where

$\frac{\partial U_{Buyer\ 2}(v,b)}{\partial b} = (v - b)(\frac{1}{5}(1 - \frac{b}{10})) - (1 - (1 - (b/10))^2)$. Solving for

$\frac{\partial U_{Buyer\ 2}(v,b)}{\partial b} = 0$ and requiring $b \leq v$ over $v \in [0, 10]$ yields:

$$b_{Buyer\ 2}(v) = \frac{1}{3}\left(v + 20 - \sqrt{v^2 - 20v + 400}\right)$$

From these two bidding functions, a simulation of 100,000,000 auctions estimates that the expected publisher revenue is 3.920.[1179]

678. *Next,* consider when Buyer Two bids first and that bid serves as the floor price for the bidders of Buyer One. Given that Buyer One runs a second-price auction, the bid function for each bidder in Buyer One is $b_{Buyer\ 1,\ Advertiser\ i}(v) = v$; the floor price does not change a bidder's optimal bid in a second-price auction. On the other hand, when Buyer Two has a valuation of *v*, it best-responds by submitting the bid *b* that maximizes its expected surplus:

$$U_{Buyer\ 2}(v,\ b) = (v - b)\ Pr\ \{Buyer\ 2\ Wins\}$$

$$= (v - b)\ (Pr\ \{\ b > v_{Buyer\ 1,\ Advertiser\ 1}\ b > v_{Buyer\ 1,\ Advertiser\ 2}\})$$

$$= (v - b)F(b)^2$$

$$= (v - b)(b/10)^2$$

[1179] This result was calculated using code/ob_example_extension.py in my supporting materials, and the output is saved in code/logs/ob_example_extension.txt.

Solving for $\frac{\partial U_{Buyer\,2}(v, b)}{\partial b} = \frac{1}{5}(v - b)\frac{b}{10} - (b/10)^2 = 0$ yields $b_{Buyer\,2}(v) = \frac{2}{3}v$. From these two bidding functions, a simulation of 100,000,000 auctions estimates that the expected publisher revenue is 4.568, a 16.5% increase over the former case.[1180]

G. Technical Notes for Section XIV

1. An Example of Post-Auction Discounts Compared To Exchange-Discriminatory Floor Prices

679. I adapt the example used by Professor Weinberg in which there is a single AdX advertiser with value V_{AdX} distributed uniformly between 12 and 16, and a single OpenX advertiser with value V_{OpenX} distributed uniformly between 10 and 12.[1181, 1182]

680. The optimal mechanism can be obtained using the standard result of Myerson (1981) which calls for allocating the impression to the bidder with the highest marginal revenue as long as it is positive.[1183] The marginal revenues are given by $MR_{AdX}(V_{AdX}) = 2V_{AdX} - 16$, and $MR_{OpenX}(V_{OpenX}) = 2V_{OpenX} - 12$.

681. Note the marginal revenues are always positive in the relevant ranges. Thus, the optimal mechanism always calls for allocating the impression. Next we compare the marginal

---

1180 This result was calculated using code/ob_example_extension.py in my supporting materials, and the output is saved in code/logs/ob_example_extension.txt.

1181 Expert Report of M. Weinberg (June 7, 2024) at ft. 248 ("To make this example mathematically rigorous, imagine that there is a single AdX bidder whose value is distributed uniformly on [12,16], and there is a single OpenX bidder whose value is distributed uniformly on [10,12]. In this case, the revenue-optimal personalized reserves are $13 on AdX and $10 on OpenX. Moreover, the example bids for AdX and OpenX are optimal in this case.").

1182 Code for computing the integration steps in this example can be found in code/post_auction_discounts_example.py in my supporting materials, and the output is saved in code/logs/post_auction_discounts_example.txt.

1183 Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58-73.

revenues to determine who to allocate the impression to. The marginal revenues are equal when $2V_{AdX} - 16 = 2V_{OpenX} - 12$, or equivalently $V_{AdX} - V_{OpenX} = 2$.

682. This implies that the optimal mechanism allocates the impression to the AdX bidder if and only if its value is at least \$2 higher than the value of the OpenX bidder and to the OpenX bidder otherwise. This allocation can be implemented with a second price auction with a \$2 post-auction discount for the OpenX bidder. In this auction, it is a dominant strategy for the AdX bidder to bid its value and for the OpenX bidder to bid its value + \$2. The expected revenue from this auction can be calculated as

$$\int_{12}^{16}\int_{10}^{12}\left(max\left\{2V_{AdX} - 16, 2V_{OpenX} - 12\right\}\frac{1}{8}\right)dV_{OpenX}dV_{AdX} = 12.333.$$

683. The expected revenue that would be attained with the two reserve prices of \$10 for OpenX and \$13 for AdX suggested by Professor Weinberg would be

$$13 * \frac{16-13}{4} + \left(1 - \frac{16-13}{4}\right) * 10 = 12.25.$$

684. Thus, publishers would be *strictly* better off using a post-auction discount of \$2 with a common reserve price of \$12 than using two different reserve prices.

685. To adapt this analysis to the context of a first-price auction (as in the UFPA) I will demonstrate that a post-auction discount of \$2 for OpenX (with a uniform reserve price of \$12 or less) continues to deliver a higher expected seller revenue than setting two reserve prices.

686. The analysis uses the derivation of optimal bidding strategies in the asymmetric first-price auction case found in Kaplan and Zamir (2012).[1184] It is convenient to start by re-writing the values as $V_{AdX} = 12 + Uniform\,[0, 4]$ and

$V_{OpenX} + \$2\,discount = 12 + Uniform\,[0, 2]$.

687. The bidding strategies can be obtained by inverting equations (20) and (21) in Kaplan and Zamir (2012),[1185] using the distributions $Uniform\,[0, 4]$ and $Uniform\,[0, 2]$ and then adding $12. Doing that, I obtain

$$b_{AdX}(V) = \frac{1}{3V}\left(4\sqrt{3V^2 + 16} - 16\right) + 12\text{, and}$$

$$b_{OpenX}(V) = \frac{2}{3V}\left(8 - 2\sqrt{16 - 3V^2}\right) + 12.$$

The expected gross revenue is then given by the following integral

$$12 + \int_0^2\int_0^4\left(max\left\{\frac{1}{3V_{AdX}}\left(4\sqrt{3V_{AdX}^{\ 2} + 16} - 16\right), \frac{2}{3V_{OpenX}}\left(8 - 2\sqrt{16 - 3V_{OpenX}^{\ 2}}\right)\right\} * \frac{1}{8}\right)dV_{AdX}dV_{OpenX},$$

which when computed, obtains $GR = 12.91801$.

688. The last step is then to subtract the post-auction discounts that need to be rebated to OpenX. For this, we compute the probability of winning for OpenX which is simply the probability that $b_{OpenX} > b_{AdX}$, which can be shown to be $\frac{1}{3}$. Thus, we subtract $\frac{1}{3} * 2$

[1184] Kaplan, T. R., & Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, *50*(2), 269-302.

[1185] Kaplan, T. R., & Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, *50*(2), 269-302.

from the gross revenue to obtain the **net revenue** $NR = GR - rebated\ amount$,

which can be calculated as $NR = 12.91801 - \frac{2}{3} = 12.25134$.

Paul R. Milgrom, Ph.D.

July 30, 2024

APPENDIX A: CURRICULUM VITAE

Paul Robert Milgrom

Shirley and Leonard Ely, Jr. Professor of Humanities and Sciences
Department of Economics
Stanford University

579 Jane Stanford Way, Landau Economics Building
Stanford, CA 94305
+1 (650) 723-3397
milgrom@stanford.edu
milgrom.people.stanford.edu

Personal

Born: April 20, 1948 in Detroit, Michigan
Spouse: Eva Meyersson Milgrom

Education

Ph.D. in Business, Stanford University, January 1979

M.S. in Statistics, Stanford University, April 1978

A.B. in Mathematics with high honors, University of Michigan, May 1970

Employment

2007–present	Senior Fellow, SIEPR, Stanford University
1993–present	Shirley and Leonard Ely, Jr. Professor of Humanities and Sciences, Stanford University
1987–present	Professor of Economics, Stanford University Professor (by courtesy), Graduate School of Business Professor (by courtesy), Department of Management Science and Engineering
1989–91	Director, Stanford Institute for Theoretical Economics

1985–87	Williams Brothers Professor of Management Studies and Professor of Economics, Yale University
1983–85	Professor of Economics and Management, Yale University
1982–83	Visiting Professor, Yale University
	Professor, Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University
1981–82	Associate Professor, Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University
1979–81	Assistant Professor, Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University

Honors, Awards, Prizes, Fellowships, and Grants

2024	Honorary Doctorate, Charles University
2023	Distinguished Research Professor at Simons Laufer Mathematical Sciences Institute (supported by the Alfred P. Sloan Foundation)
2020	Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel
	Distinguished Fellow of the American Economics Association
2019	Robert Rosenthal Memorial Lecture (Boston University),
	Aumann Lecture (Israeli Game Theory Society),
	Marshall Lectures (Cambridge University)
2018	John J. Carty Award for the Advancement of Science, U.S. National Academy of Sciences
2017	CME Group-MSRI prize in Innovative Quantitative Applications, Chicago Mercantile Exchange and Mathematical Sciences Research Institute
	McKenzie lecture, University of Rochester
	Stanford Humanities and Sciences Dean's Award for Excellence in Graduate Education
	Elected Fellow of the Game Theory Society
	Elected Fellow of the Finance Theory Group
2016	Nancy Schwartz Memorial Lecture, Northwestern University

2015	National Science Foundation Award "Auction Market Design"
	Simon's Institute Public Lecture, University of California, Berkeley
	WINE (Web and Internet Economics) Keynote Lecture
2014	Golden Goose Award
	Keyfitz Lecture, Fields Institute, Toronto
	Arrow Lecture, Columbia University
2013	Nomura Lecturer, Institute of Mathematics, Oxford University
	BBVA Foundation Frontiers of Knowledge Award in Economics, Finance and Management
2012	Elected Vice President of the American Economic Association (term to begin in 2013)
	Inaugural lecture on "Incentive Auctions for Radio Spectrum," C.V. Starr Center Distinguished Speaker Series, New York University
	Oskar Morgenstern lecture on "Designing the US 'Incentive Auctions,'" Fourth World Congress of the Game Theory Society
	Becker Friedman Visitor, University of Chicago
	Intertic Stackelberg Lecture on "Auctions for Online Display Advertising"
2011	Eitan Berglas Lecture on "The Applied Science of Market Design," Tel Aviv University
2010	NSF-SBIR Phase IB Award for "Incorporating Bidder Budget Constraints in Multi-item Auctions"
2009	NSF-SBIR Phase I Award for "Incorporating Bidder Budget Constraints in Multi-Item Auctions"
	Nemmers Lecture, Northwestern University
	EARIE (European Association for Research in Industrial Organization) Lecture
2008	Erwin Plein Nemmers Prize
	W.A. Mackintosh Lecture, Queens University
	Simon Newcomb Lecture, Johns Hopkins University
2007	President, Western Economic Association International (WEAI)
	National Science Foundation Grant on "Market Design"

2006	Elected to the National Academy of Sciences
	Colin Clarke Lecture, Econometric Society Australasian Meeting
	Manchot Lecture, University of Bonn
2005	Elected to the Executive Committee of the Econometric Society
	Elected Vice President of the Western Economic Association
	Chung-Hua Lecturer, Academia Sinica (Taiwan)
	Clarendon Lecturer, Oxford University
2004	Fischer-Schulz Lecturer, Econometric Society
	Koopmans Lecturer, Yale University
	National Science Foundation Research Grant to study "Electronic Auction Markets"
	Council Member, Econometric Society
2003	National Science Foundation Research Grant to study "Cumulative Offer Processes"
	Landau Economics Teaching Prize, Stanford University
	Elected to the Council, Game Theory Society
	Distinguished Economist Lecture, Federal Communications Commission
2001	Honorary Doctorate, Stockholm School of Economics
2000	Taussig Visiting Research Professor, Harvard University
1999	Murray S. Johnson Inaugural Lecture, University of Texas
	Industry Canada Distinguished Lecture
1998	Fain Lecture, Brown University
	Lawrence Klein Lecture, University of Pennsylvania
	Fellow (2nd time), Center for Advanced Study in the Behavioral Sciences
1997	Alberto Bailleres Founder's Lecture at Instituto Tecnologico Autonomo de Mexico (ITAM)
	Plenary Lecturer, Econometric Society Far Eastern Meeting
	Plenary Lecturer, Australian Industry Economics Meeting, University of Melbourne
1996	Nobel Prize Memorial Lecture (honoring deceased Nobel laureate William Vickrey) at the Royal Swedish Academy of Sciences
1995	Churchill Lectures at Cambridge University
	Political Economy Special Lecture at Harvard University

1994	National Science Foundation Research Grant to study "Comparative Statics, Complementarities, Coordination and Change," (covering 1994 to 1997)
	Woytinsky Distinguished Lecturer, University of Michigan
1993	Senior Research Fellow, Institute for Policy Reform
	Shirley R. and Leonard W. Ely, Jr. Professor of Humanities and Sciences, Stanford University
1992	Fellow, American Academy of Arts and Sciences
	International Guest Scholar, Kyoto University
1991	Fellow, Center for Advanced Study in the Behavioral Sciences
	National Science Foundation Research Grant to study "Theories of the Firm 2" (covering 1991 to 1994)
1990	Center for Economic Policy Research Grant to study "The Economics of Modern Manufacturing"
1989	National Science Foundation Grant to direct programs for the Stanford Institute for Theoretical Economics
	National Academy of Sciences Award to lecture in China on economics of organizations
1988	Olin Distinguished Lecturer, Princeton University
	National Science Foundation Research Grant to study "Theories of the Firm" (covering 1988 to 1991)
	Center for Economic Policy Research Grant
1987	Prize for Best Paper of the Year in the Transactions of the Society of Actuaries
1986	Ford Visiting Professor of Economics, University of California, Berkeley
	John Simon Guggenheim Fellowship to study "Economic Theories of Organization"
1985	Williams Brothers Chair in Management Studies, Yale University
	National Science Foundation Research Grant to study "On the Formal Economic Theory of Organizations"
	Fellow of the Institute for Advanced Studies, Hebrew University of Jerusalem
	Plenary Lecturer at the Fifth World Congress of the Econometric Society

1984	Fellow of the Econometric Society
	Fellow of Morse College, Yale University
1983	Research Award, Actuarial Education and Research Fund
	Honorary Master of Arts degree, Yale University
1982	National Science Foundation Research Grant to study "The Structure of Information in a Productive Organization"
1981	IBM Research Chair at Northwestern University
	Visiting Research Associate, Stanford University
1980	Leonard J. Savage Memorial Thesis Award
	National Science Foundation Research Grant to study "Information and Uncertainty in Competitive Bidding"
1976	Society of Actuaries Triennial Paper Prize for best paper by an actuary within five of membership, for the period 1973–75
1974	Fellow of the Society of Actuaries

Publications

Articles

"Incentive Auction Design Alternatives: A Simulation Study," with Kevin Leyton-Brown, Neil Newman and Ilya Segal. *Management Science,* February 2024.

"Algorithmic Mechanism Design with Investment," with Mohammad Akbarpour, Scott Kominers, Kevin Michael Li, and Shengwu Li. *Econometrica,* November 2023.

"Taming the Communication and Computation Complexity of Combinatorial Auctions: The FUEL Bid Language," with Martin Bichler and Gregor Schwarz. *Management Science*, June 2022.

"When Should Control Be Shared?" with Eva Meyersson Milgrom and Ravi Singh, *Management Science*, March 2022.

"Auction Research Evolving: Theorems and Market Designs," *American Economic Review,* May 2021.

"Extended Proper Equilibrium," with Joshua Mollner, *Journal of Economic Theory*, June 2021.

“Clock Auctions and Radio Spectrum Reallocation,” with Ilya Segal, *Journal of Political Economy*, January 2020.

“Auction Market Design: Recent Innovations,” *Annual Reviews*, August 2019.

“Equilibrium Selection in Auctions and High Stakes Games,” with Joshua Mollner, *Econometrica*, January 2018.

“Redesigning Spectrum Licenses,” with Anthony Zhang and E. Glen Weyl, *Regulation*, Fall 2017.

“Economics and Computer Science of a Radio Spectrum Reallocation,” with Kevin Leyton-Brown and Ilya Segal, *Proceedings of the National Academy of Sciences*, July 2017.

“Adverse Selection and Auction Design in Internet Display Advertising,” with Nick Arnosti and Marissa Beck, *American Economic Review*, October 2016.

“Ascending Prices and Package Bidding: Further Experimental Analysis,” with John Kagel and Yuanchuan Lien, *Games and Economic Behavior*, May 2014.

“Designing Random Allocation Mechanisms: Theory and Applications,” with Eric Budish, Yeon-Koo Che and Fuhito Kojima, *American Economic Review*, April 2013.

“Critical Issues in Market Design,” *Economic Inquiry*, April 2011.

“Simplified Mechanisms with an Application to Sponsored-Search Auctions,” *Games and Economic Behavior*, September 2010.

“Ascending Prices and Package Bidding: A Theoretical and Experimental Analysis,” with Yuanchuan Lien and John Kagel, *American Economic Journal: Microeconomics*, August 2010.

“Online Advertising: Heterogeneity and Conflation in Market Design,” with Jon Levin, *American Economic Review*, May 2010.

“Assignment Messages and Exchanges,” *American Economic Journal: Microeconomics*, August 2009. Reprinted in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

“The Limited Influence of Unemployment on the Wage Bargain,” with Robert Hall, *American Economic Review*, September 2008. Reprinted in *Political Economy: Critical Concepts*, by Norman Schofield, Dino Falaschetti and Andrew Rutten (eds.), New York: Routledge, 2011.

“Substitute Goods, Auctions and Equilibrium,” with Bruno Strulovici, Journal *of Economic Theory*, June 2008

“The Promise of Prediction Markets” (with 22 co-authors), *Science*, May 2008.

“What the Seller Won’t Tell You: Persuasion and Disclosure in Markets,” *Journal of Economics Perspectives*, Spring 2008.

“Core-Selecting Package Auctions,” with Bob Day, *International Journal of Game Theory*, March 2008. Reprinted in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

“Package Auctions and Package Exchanges” (2004 Fisher-Schultz lecture), *Econometrica*, July 2007.

“Matching with Contracts,” with John Hatfield, *American Economic Review*, September 2005.

“Ascending Auctions with Package Bidding,” with Lawrence M. Ausubel, *Frontiers of Theoretical Economics*, August 2002. Republished in *Advances in Theoretical Economics*, 2002.

“Package Bidding: Vickrey vs Ascending Auctions,” with Lawrence M. Ausubel, *Revue Economique*, May 2002.

“Envelope Theorems for Arbitrary Choice Sets,” with Ilya Segal, *Econometrica*, March 2002.

“Advances in Routing Technologies and Internet Peering Agreements,” with Stan Besen, Bridger Mitchell and Padmanabhan Srinagesh, *American Economic Association Papers and Proceedings*, May 2001.

“Putting Auction Theory to Work: The Simultaneous Ascending Auction,” *Journal of Political Economy*, April 2000. Reprinted in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

“Game Theory and the Spectrum Auctions,” *European Economic Review*, May 1998.

“Coalition-Proofness and Correlation with Arbitrary Communication Possibilities,” with John Roberts, *Games and Economic Behavior*, November 1996.

“The LeChatelier Principle,” with John Roberts, *American Economic Review*, March 1996. Reprinted in *Paul Anthony Samuelson, Critical Assessments of Contemporary Economists*, by John Cunningham Wood and Michael McLure (eds.), New York: Routledge, 2004.

“The Economics of Modern Manufacturing: Technology, Strategy and Organization: Reply,” *American Economic Review*, September 1995.

"Deterring Predation in Telecommunications: Are Line-of-Business Restraints Needed?" with Susan Gates and John Roberts, *Managerial and Decision Economics*, July 1995. Reprinted in *Deregulating Telecommunications: The Baby Bells' Case for Competition*, by R.S. Higgins and P.H. Rubin (eds). Chichester: John Wiley & Sons, 1995.

"Complementarities and Fit: Strategy, Structure and Organizational Change in Manufacturing," with John Roberts, *Journal of Accounting and Economics*, March 1995.

"The Firm as an Incentive System," with Bengt Holmstrom, *American Economic Review*, September 1994. Reprinted in *The Theory of the Firm: Critical Perspectives*, by Nicolai Juul Foss (ed.), New York: Routledge, 2000. Reprinted in *Readings in the Economics of the Division of Labor, Vol 2: Modern Analyses*, by Guang-Zhen Sun (ed.), World Scientific, 2005.

"Coordination, Commitment and Enforcement: The Case of the Merchant Guild," with Avner Greif and Barry Weingast, *Journal of Political Economy*, August 1994. Reprinted in *Explaining Social Institutions*, by Jack Knight and Itai Sened (eds.), Ann Arbor: University of Michigan Press, 1995. Reprinted in *Trust*, by Elias Khalil (ed.), London: Edward Elgar Publishing, 2002. Reprinted in *The Foundations Library of the New Institutional Economics*, by Claude Ménard (ed.), London: Edward Elgar Publishing, 2005. Reprinted in *Social Norms, Non-legal Sanctions, and the Law*, by Eric Posner (ed.), London: Edward Elgar Publishing, 2007. Reprinted in *Customary Law*, by Lisa Bernstein and Francesco Parisi (eds.), London: Edward Elgar Publishing, forthcoming.

"Comparing Equilibria," with John Roberts, *American Economic Review*, June 1994. Reprinted in *Equilibrium*, by Donald Walker (ed.), Edward Elgar Publishing, 2000.

"Comparing Optima: Do Simplifying Assumptions Affect Conclusions?" *Journal of Political Economy*, June 1994.

"Complementarities and Systems: Understanding Japanese Economic Organization," with John Roberts, *Estudios Economicos*, April 1994.

"Monotone Comparative Statics," with Chris Shannon, *Econometrica*, January 1994.

"Organizational Prospects, Influence Costs and Ownership Changes," with Margaret Meyer and John Roberts, *Journal of Economics and Management Strategy*, Spring 1992.

"Pay, Perks and Parachutes: Do They Pay?" with John Roberts, *Stanford Business*, 1992.

"Information and Timing in Repeated Partnerships," with Dilip Abreu and David Pearce, *Econometrica*, November 1991.

"A Theory of Hierarchies Based on Limited Managerial Attention," with John Geanakoplos, *Journal of Japanese and International Economies*, September 1991. Reprinted in *The Economics of Organization and Bureaucracy*, by Peter Jackson (ed.), London: Edward Elgar Publishing, 2013.

"Complementarities, Momentum, and the Evolution of Modern Manufacturing," with Yingyi Qian and John Roberts, *American Economic Association Papers and Proceedings*, May 1991.

"Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design," with Bengt Holmstrom, *Journal of Law, Economics and Organization*, Spring 1991. Reprinted in *Transaction Cost Economics*, by Oliver Williamson and Scott Masten (eds.), London: Edward Elgar Publishing, 1994. Reprinted in *The Principal-Agent Model: The Economic Theory of Incentives*, by J-J Laffont (ed.), Cheltenham: Edward Elgar Press, 2003. Reprinted in *The International Library of the New Institutional Economics*, by Claude Ménard (ed.), London: Edward Elgar Publishing, 2005. Reprinted in *The Economic Nature of the Firm*, by Louis Putterman and Randall Kroszner (ed.), Cambridge University Press, 1996. Reprinted in *The Economics of Contracts*, by Patrick Bolton, Barbara and David Zalaznick (eds.), Cheltenham: Edward Elgar Press, 2008. Reprinted in *Institutional Law and Economics*, by Pablo Spiller (ed.), Cheltenham: Edward Elgar Press, forthcoming.

"Adaptive and Sophisticated Learning in Repeated Normal Form Games," with John Roberts, *Games and Economic Behavior*, February 1991. Reprinted in *Recent Developments in Game Theory*, by E. Maskin (ed.), Cheltenham: Edward Elgar, 1998.

"Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities," with John Roberts, *Econometrica*, November 1990. Reprinted in *Recent Developments in Game Theory*, by E. Maskin (ed.), Cheltenham: Edward Elgar, 1998.

"The Economics of Modern Manufacturing: Technology, Strategy and Organization," with John Roberts, *American Economic Review*, June 1990. Reprinted in *Vestnik St. Petersburgskogo Universiteta, Economics Seria* (the Journal of the Economics Faculty of St. Petersburg University), 1993 (in translation). Reprinted in *Economics of the Firm: Lessons in Business Organization*, by Andrei Demin and Valery Katkalo (eds.), St. Petersburg, Russia: 1994. Reprinted in *The Economics of Communications and Information*, by Donald Lamberton (ed.), Cheltenham: Edgar Elgar Publishing, 1996. Reprinted in *Readings in Applied Microeconomic Theory: Market Forces and Solutions*, by Robert E. Kuenne (ed.), Blackwell Publishers, 2000. Reprinted in *Fundamentals of Business Strategy*, by Mie Augur and David Teece (eds.), Sage Publications, 2007.

"Short Term Contracts and Long Term Agency Relationships," with Drew Fudenberg and Bengt Holmstrom, *Journal of Economic Theory*, June 1990.

"The Efficiency of Equity in Organizational Decision Processes," with John Roberts, *American Economic Review Papers and Proceedings*, May 1990.

"The Role of Institutions in the Revival of Trade: The Medieval Law Merchant," with Douglass North and Barry Weingast, *Economics and Politics*, March 1990. Reprinted in *Trade in the Pre-Modern Period: 1400-1700*, by Douglas Irwin (ed.), London: Edward Elgar Publishing, 1996. Reprinted in *Reputation: Studies in the Voluntary Elicitation of Good Conduct*, by Daniel B. Klein (ed.), Ann Arbor, University of Michigan Press, 1997. Reprinted in *The Political Economy of Institutions*, by Claude Ménard (ed.), London: Edward Elgar Publishing, 2004. Reprinted in *International Institutions in the New Global Economy*, by Lisa L. Martin (ed.), London: Edward Elgar Publishing, 2005. Reprinted in *Anarchy and the Law*, by Edward Stringham (ed.), New Brunswick, New Jersey: Transaction Publishers, 2006. Reprinted in *Social Norms, Non-Legal Sanctions, and the Law*, by Eric A. Posner (ed.), London: Edward Elgar Publishing, 2007.

"Regulating Trade Among Agents," with Bengt Holmstrom, *Journal of Institutional and Theoretical Economics*, March 1990. Reprinted in *The New Institutional Economics*, by Erik G. Furubotn and Rudolph Richter (eds.), College Station: Texas A&M University Press, 1991.

"Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, Summer 1989. Reprinted in *Readings in Microeconomic Theory*, by Manfredi La Manna (ed.), London: Dryden Press, 1997.

"Communication and Inventories as Substitutes in Organizing Production," with John Roberts, *Scandinavian Journal of Economics*, September 1988.

"Economic Theories of Organization: Past, Present and Future," with John Roberts, *Canadian Journal of Economics*, August 1988. Reprinted in *The Economics of Contracts and Industrial Organization: A Reader*, Peter Buckley and Jonathan Michie (eds.), Oxford University Press, 1996.

"An Economic Approach to Influence Activities and Organizational Responses," with John Roberts, *American Journal of Sociology*, July 1988.

"Employment Contracts, Influence Activities and Efficient Organization Design," *Journal of Political Economy*, February 1988.

"Job Discrimination, Market Forces, and the Invisibility Hypothesis," with Sharon Oster, *Quarterly Journal of Economics*, August 1987. Reprinted in *Learning in Labour Markets*, by Michael Waldman (ed.), Cheltenham: Edward Elgar Press.

"Informational Asymmetries, Strategic Behavior and Industrial Organization," with John Roberts, *American Economic Association Papers and Proceedings*, May 1987.

"Aggregation and Linearity in the Provision of Intertemporal Incentives," with Bengt Holmstrom, *Econometrica*, March 1987. Reprinted in *The Principal-Agent Model: The Economic Theory of Incentives*, by J-J Laffont (ed.), Cheltenham: Edward Elgar Press. Reprinted in *The Economics of Contracts*, by Patrick Bolton, Barbara and David Zalaznick (eds.), Cheltenham: Edward Elgar Press.

"Price and Advertising Signals of Product Quality," with John Roberts, *Journal of Political Economy*, August 1986. Reprinted in *Antitrust and Competition*, by Andrew Kleit (ed.), Cheltenham: Edward Elgar Publishing, 2005. Reprinted in *The Economics of Marketing, Martin Carter*, by Mark Casson and Vivek Suneja (eds.), Cheltenham: Edgar Elgar Publishing, 1998. Reprinted in *Readings in Industrial Organization*, by Luis M.B. Cabral (ed.), Oxford: Blackwell Publishers, 2000.

"Relying on the Information of Interested Parties," with John Roberts, *Rand Journal of Economics*, April 1986. Reprinted in *Economics of Evidence, Procedure and Litigation*, by Chris William Sanchirico (ed.), Cheltenham: Edward Elgar Publishing, 2007.

"Distributional Strategies for Games with Incomplete Information," with Robert Weber, *Mathematics of Operations Research*, November 1985.

"Bid, Ask and Transactions Prices in a Specialist Market with Heterogeneously Informed Traders," with Lawrence R. Glosten, *Journal of Financial Economics*, March 1985.

“Reply to the Comments on ‘Measuring the Interest Rate Risk,’” *Transactions of the Society of Actuaries, XXXVII*, 1985

“Measuring the Interest-Rate Risk,” *Transactions of the Society of Actuaries, XXXVII*, 1985.

“Competitive Bidding with Proprietary Information,” with Richard Engelbrecht-Wiggans and Robert Weber, *Journal of Mathematical Economics*, April 1983.

“A Theory of Auctions and Competitive Bidding,” with Robert Weber, *Econometrica*, September 1982. Reprinted in *Game Theory in Economics*, by Ariel Rubinstein (ed.), London: Edward Elgar Publishing, 1990. Reprinted in *The Economics of Information*, by Steven A. Lippman and John E. Anderson (eds.), London: Edward Elgar Publishing, 1994. Reprinted in *The Economic Theory of Auctions*, by Paul Klemperer (ed.), London: Edward Elgar Publishing, 1999. Reprinted in *Readings in Applied Microeconomic Theory: Market Forces and Solutions*, by Robert E. Kuenne (ed.), Blackwell Publishers, 2000.

“Predation, Reputation, and Entry Deterrence,” with John Roberts, *Journal of Economic Theory*, August 1982. Reprinted in *Antitrust and Competition*, by Andrew Kleit (ed.), Cheltenham: Edward Elgar Publishing, 2005. Reprinted in *The Economics of Reputation*, by Jill J. McCluskey and Jason Winfree (eds.), Edward Elgar Publishing, 2017.

“Rational Cooperation in the Finitely-Repeated Prisoners’ Dilemma,” with David Kreps, John Roberts and Robert Wilson, *Journal of Economic Theory*, August 1982. Reprinted in *Game Theory in Economics*, by Ariel Rubinstein (ed.), Cheltenham: Edward Elgar Publishing, 1990. Reprinted in *Trust*, by Elias L. Khalil (ed.), Cheltenham: Edward Elgar Publishing, 2003.

“The Value of Information in a Sealed Bid Auction,” with Robert Weber, *Journal of Mathematical Economics*, June 1982. Reprinted in *The Economic Theory of Auctions*, by Paul Klemperer (ed.), London: Edward Elgar Publishing, 1999.

“Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis,” with John Roberts, *Econometrica*, March 1982. Reprinted in *Industrial Organization*, by Oliver Williamson (ed.), London: Edward Elgar Publishing, 1990. Reprinted in *The Economics of Information*, by Steven A. Lippman and John E. Anderson (eds.), London: Edward Elgar Publishing, 1994. Reprinted in *Readings in Industrial Organization*, edited by Luis M.B. Cabral (ed.), Oxford: Blackwell Publishers, 2000. Reprinted in *Pricing*, by Michael Waldman and Justin P. Johnson (ed.), London: Edward Elgar Publishing, 2007.

"Information, Trade and Common Knowledge," with Nancy Stokey, *Journal of Economic Theory*, February 1982.

"Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, October 1981.

"Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, July 1981. Reprinted in *The Economic Theory of Auctions*, by Paul Klemperer (ed.), London: Edward Elgar Publishing, 1999.

"An Axiomatic Characterization of Common Knowledge," *Econometrica*, January 1981.

"A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica*, May 1979.

"On Understanding the Effects of GAAP Reserve Assumptions," *Transactions of the Society of Actuaries, XXVII*, 1975.

Books

Discovering Prices: Auction Design in Markets with Complex Constraints. Columbia University Press, 2017.

Putting Auction Theory to Work. Cambridge: Cambridge University Press, 2004.

Contributions to Books and Proceedings

"Spectrum Auctions from the Perspective of Matching," with Andrew Vogt, *Online and Matching-Based Market Design*, Cambridge University Press, 2023.

"How Artificial Intelligence and Machine Learning Can Impact Market Design," with Steve Tadelis, *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay K. Agrawal, Joshua Gans, and Avi Goldfarb, 2019.

"Winning Play in Spectrum Auctions," with Jeremy Bulow and Jonathan Levin, in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

"Designing the US Incentive Auction," in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

"Optimal Incentives in Core-Selecting Auctions," with Bob Day, in *The Handbook of Market Design*, by Zvika Neeman, Al Roth, and Nir Vulkan (eds.), Oxford University Press, 2013.

“Incentive Auction: Rules and Discussion,” with Lawrence Ausubel, Jonathan Levin and Ilya Segal, published as Appendix C of the FCC 12-118 (Notice of Proposed Rulemaking on the Incentive Auction, Released October 2, 2012).

“Complementarity in Organizations,” with Erik Brynjolfsson, in *The Handbook of Organizational Economics*, by John Roberts and Bob Gibbons (eds.), Princeton University Press, 2012.

“Multipliers and the LeChatelier Principle,” in *Samuelsonian Economics and the Twenty-First Century*, by Michael Szenberg, Lall Ramrattan and Aron Gottesman (eds.), Oxford University Press, 2006.

“The Clock-Proxy Auction: A Practical Combinatorial Design,” with Lawrence M. Ausubel and Peter Cramton, in *Combinatorial Auctions*, by Peter Cramton, Richard Steinberg and Yoav Shoham (eds.), MIT Press, 2005. Reprinted in *Handbook of Spectrum Auction Design*, by Martin Bichler and Jacob Goeree (eds.), Cambridge University Press, 2017.

“Ascending Proxy Auctions,” with Lawrence M. Ausubel, in *Combinatorial Auctions*, by Peter Cramton, Richard Steinberg and Yoav Shoham (eds.), MIT Press, 2005.

“The Lovely but Lonely Vickrey Auction,” with Lawrence M. Ausubel, in *Combinatorial Auctions*, by Peter Cramton, Richard Steinberg and Yoav Shoham (eds.), MIT Press, 2005.

“Competitive Effects of Internet Peering Policies,” with Bridger Mitchell and Padmanabhan Srinagesh, in *The Internet Upheaval*, by Ingo Vogelsang and Benjamin Compaine (eds.), Cambridge: MIT Press, 2000.

“A Theory of Auctions and Competitive Bidding, II,” with Robert Weber, in *The Economic Theory of Auctions*, by Paul Klemperer (ed.), Edward Elgar Publishing, 1999.

“Combination Bidding in Spectrum Auctions,” in *Competition, Regulation and Convergence: Current Trends in Telecommunications Research*, by Sharon Gillett and Ingo Vogelsang (eds.), New Jersey: Lawrence Erlbaum Associates, Publishers, 1999.

“The Internal Politics of the Firm,” with John Roberts, in *The Politics of Exchange and the Economics of Power*, by Samuel Bowles, Maurizio Franzini and Ugo Pagano (eds.), New York: Routledge, 1998.

“Procuring Universal Service: Putting Auction Theory to Work,” in *Le Prix Nobel: The Nobel Prizes, 1996*, Nobel Foundation, 1997.

“Procuring Universal Telephone Service,” in *1997 Industry Economics Conference, Industry Commission (ed.), Conference Proceedings*, 1997, AGPS.

"Complementarities in the Transition from Socialism: A Firm-Level Analysis," with Susan Gates and John Roberts, in *Reforming Asian Socialism: The Growth of Market Institutions*, by John McMillan and Barry Naughton (eds.), Ann Arbor: University of Michigan Press, 1996.

"Continuous Adjustment and Fundamental Change in Business Strategy and Organization," with John Roberts, in *Trends in Business Organization: Do Participation and Cooperation Increase Competitiveness?*, by Horst Siebert (ed.), Tübingen: J.C.B. Mohr (Paul Siebeck), 1995.

"Johnson Controls, Inc—Automotive Systems Group: The Georgetown Kentucky Plant," with John Roberts, *Case #S-BE-9, Stanford Graduate School of Business*, November 1993.

"Is Sympathy an Economic Value? Philosophy, Economics and the Contingent Valuation Method," in *Contingent Valuation: A Critical Assessment*, by Jerry Hausman (ed.), Elsevier-North Holland, 1993.

"The Real Output of the Stock Exchange," with Timothy F. Bresnahan and Jonathan Paul, in *Output Measurement in the Services Sectors*, by Zvi Griliches (ed.), 1992.

"Bargaining Costs, Influence Costs and the Organization of Economic Activity," with John Roberts, in *Perspectives on Positive Political Economy*, by James E. Alt and Kenneth A. Shepsle (eds.), Cambridge: Cambridge University Press, 1990. Reprinted in *Transaction Cost Economics*, by Oliver Williamson and Scott Masten (eds.), London: Edward Elgar Publishing Co., 1994. Reprinted in *The International Library of the New Institutional Economics*, by Claude Ménard (ed.), London: Edward Elgar Publishing, 2003. Reprinted in *The Economic Nature of the Firm*, by Louis Putterman and Randall Kroszner (eds.), Cambridge University Press, 1996. Reprinted in *The International Library of the New Institutional Economics*, by Claude Ménard (ed.), London: Edward Elgar Publishing, 2005.

"New Theories of Predatory Pricing," with John Roberts, *Industrial Structure in the New Industrial Economics*, by Giacomo Bonanno and Dario Brandolini (eds.), Oxford: Oxford University Press, 1990.

"An Essay on Price Discrimination," in *The Economics of Imperfect Competition and Employment: Joan Robinson and Beyond*, by George Feiwel (ed.), New York: MacMillan and New York: New York University Press, 1989.

"Predatory Pricing," in *The New Palgrave: A Dictionary of Economic Theory and Doctrine*, by J. Eatwell, M. Milgate, and P. Newman (eds.), London: MacMillan Press Ltd., 1988.

"Auction Theory," in *Advances in Economic Theory: Fifth World Congress*, by Truman Bewley (ed.), London: Cambridge University Press, 1987.

"The Economics of Competitive Bidding: A Selective Survey," in *Social Goals and Social Organization: A Volume in Honor of Elisha Pazner*, by L. Hurwicz, D. Schmeidler and H. Sonnenschein (eds.), London: Cambridge University Press, 1985.

"Private Information in an Auction-Like Securities Market," in *Auctions, Bidding and Contracting: Uses and Theory*, by R. Engelbrecht-Wiggans, M. Shubik and R. Stark (eds.), New York: New York University Press, New York, 1983.

"Topologies on Information and Strategies in Normal-Form Games with Incomplete Information," with Robert Weber, in *Game Theory and Mathematical Economics*, by O. Moeschlin and D. Pallaschke (eds.), New York: North Holland, 1981.

Patents

Paul Milgrom, Stephan Cunningham, Marissa Beck. "United States Patent 11,574,358 Impression Allocation System and Method Using an Auction that Considers Losing Bids," OpenX Technologies, Inc, Feb 7, 2023

Paul Milgrom, Steve Goldband. "United States Patent 8,788,364 System for Configuration and Implementation of an Assignment Auction or Exchange," Auctionomics, Inc, Jul 22, 2014

Lawrence Ausubel, Peter Cramton, Paul Milgrom. "United States Patent 8,744,924 System and Method for Efficient Clearing of Spectrum Encumbrances," Efficient Auctions, LLC, Jun 3, 2014

Marc Porat, Kevin Surace, Paul Milgrom. "United States Patent 8,738,463 Method, System and Business Model for a Buyer's Auction with Near-Perfect Information Using the Internet," Perfect Commerce, LLC, May 27, 2014

Paul Milgrom, David Salant. "United States Patent 8,577,746 System and Method for a Multi-Product Clock Auction," Auction Technologies, LLC, Nov 5, 2013

Lawrence Ausubel, Paul Milgrom. "United States Patent 8,566,211 System and Method for a Dynamic Auction with Package Bidding," Efficient Auctions, LLC, Oct 22, 2013

Marc Porat, Kevin Surace, Paul Milgrom. "United States Patent 8,341,033 Method, System and Business Model for a Buyer's Auction with Near-Perfect Information Using the Internet," Perfect Commerce, Dec 25, 2012

Lawrence Ausubel, Peter Cramton, Paul Milgrom. "United States Patent 8,335,738 System and Method for a Hybrid Clock and Proxy Auction," Efficient Auctions, LLC, Dec 18, 2012

Paul Milgrom, Steven Goldband. "United States Patent 8,271,345 System and Method for Incorporating Bidder Budgets in Multi-Item Auctions," Auctionomics, Inc, Sep 18, 2012

Lawrence Ausubel, Peter Cramton, Paul Milgrom. "United States Patent 8,224,743 System and Method for a Hybrid Clock and Proxy Auction," Efficient Auctions, LLC, Jul 17, 2012

Others

"The Market Design Community and the Broadcast Incentive Auction: Fact-Checking Glen Weyl's and Stefano Feltri's False Claims," Digitopoly, June 2020.

"Commentary on Effective Allocation of Affordable Housing by Nick Arnosti and Peng Shi" with Mitchell Watt, Management Science Blog, May 2020.

"How Obscure Science Led to Spectrum Auctions That Connected the World," April 2017.

"Ad Exchanges Are Not All The Same: How Better Policed Exchanges Help Premium Publishers," OpenX, 2015.

"The Case for Unlicensed Spectrum," with Jonathan Levin, October 2011.

"Making Carbon Markets Work," SIEPR, May 2009.

"Using Procurement Auctions to Allocate Broadband Stimulus Grants," SIEPR, May 2009.

"The Promise of Prediction Markets," with multiple co-authors, May 2008.

"Economists' Statement on US Broadband Policy," with multiple co-authors, March 2006.

"Promoting Efficient Use of Spectrum Through Elimination of Barriers to Secondary Markets," with multiple co-authors, February 2001.

"An Economist's Vision of the B-to-B Marketplace," Palo Alto: Perfect, 2000.

Working Papers

"Kenneth Arrow's Last Theorem," May 2024.

"Walrasian Mechanisms for Non-Convex Economies and the Bound-Form First Welfare Theorem," with Mitchell Watt, July 2022.

"Fuel for 5g," Auctionomics White Paper, June, 2019.

"The CAF Auction: Design Proposal," with Assaf Eilat, July 2011.

"Adverse Selection without Hidden Information," June 1987.

Major Professional Activities and Affiliations

2017-21	Chair, Economics section (54) of the National Academy of Sciences
2020	SIGecom Test of Time Award Committee
2019-20	CME-MSRI Award Committee
2016–17	National Academy of Sciences: Class Membership Committee
	Chair, NAS Temporary Nominating Group
2016	National Academy of Sciences: Air Force Studies Board Committee
2015–present	Editorial Board, Proceedings of the National Academy of Sciences
	Executive Supervisory Committee, CERGE-EI
	National Academies' Intelligence Science and Technology Experts Group (ISTEG)
2012–17	Lead consultant to Federal Communications Commission Incentive Auctions Task Force
2012–14	Editorial Board of European Journal of Pure and Applied Mathematics
2009–present	Editorial Board of AEJ-Microeconomics
2007–08	President, Western Economic Association International (WEAI)
2006–07	Member, National Academy of Sciences
	President-Elect, Western Economic Association International (WEAI)
2005–06	Vice President, Western Economic Association International (WEAI)
2005–08	Executive Committee of the Econometric Society
2004–06	Council, Econometric Society
2003–present	Council, Game Theory Society

2000–02	Chief economist, Perfect Commerce
1997–02	Editorial Consultant, MIT Press
1997–99	Editorial Board, Journal of Comparative Economics
1996–16	Founder and Director, Market Design Inc. (Chairman, 1996–02)
1996	Nemmers Prize Selection Committee, Northwestern University
1996–06	Advisory Board, Microeconomics Abstracts
1995–05	Advisory Board, Economics Research Network
1994–95	Program Committee, 1995 World Congress of the Econometric Society
1993–95	Senior Research Fellow, Institute for Policy Reform
1993–00	Associate Editor, American Economic Review
1992–present	Fellow, American Academy of Arts and Sciences
1990–93	Co-Editor, American Economic Review
1990–present	Associate Editor, Games and Economic Behavior
1989–92	Associate Editor, Journal of Financial Intermediation
1987–90	Associate Editor, Econometrica
1985–89	Associate Editor, Rand Journal of Economics
1983–87	Associate Editor, Journal of Economic Theory
1984	Chair, Program Committee, Econometric Society Winter Meetings
1984–present	Fellow, Econometric Society
1980–present	Member, American Economic Association

Expert Reports and Testimony (2000–present)

1. Public Utility Commission of Oregon, in the matter of the Application of Portland General Electric for Approval of the Customer Choice Plan, Docket UE 102.

2. American Arbitration Association, in the matter of MCG PCS v Leap Wireless International. (February 2002)

3. American Arbitration Association, in the matter of Guilherme Augusto Frering and Mario Augusto Frering v Mitsui and Co., Ltd., and Cayman Iron Ore Investment Co., Ltd. (July 2003)

4. United States of America Ex Rel. R.C. Taylor III, v Mario Gabelli, Lynch Corp., Lynch Interactive Corp., Lynch PCS Communications Corp. et al. Case No. 03 Civ. 8762 (SAS)(GWG), Southern District of New York. (March 2006)

5. Bid For Position v. Google et al, Case No 2:07-cv- 582 (JBF/TEM), Eastern District of Virginia, Norfolk Division.

6. Marla Tidenberg v. BidZ, Inc, United States District Court, Central District of California, Case No. CV08-05553 PSG (FMOx). (January 2010).

7. Alaska Electrical Pension Fund, et al. v. Bank of America Corp., et al, Civil Action No. 14-cv-7126 (JMF). (January 2018)

8. Rick Woods v. Google LLC, Case 5:11-cv-01263-EJD. United States District Court, Northern District of California, San Jose Division.

9. United States, et al. v. Google LLC, Case 1:23-cv-0108 (LMB/JFA), United States District Court for the Eastern District of Virginia (January 2024).

APPENDIX B: LIST OF MATERIAL RELIED ON

A. *Complaint and Expert Reports*

State of Texas et al. v. Google LLC, Fourth Amended Complaint (May 5, 2023).
Expert Report of J. Andrien (Jun. 7, 2024).
Expert Report of J. Chandler (Jun. 7, 2024).
Expert Report of J. Gans (Jun. 7, 2024).
Expert Report of J. Hochstetler (Jun. 7, 2024).
Expert Report of P. Pathak (Jun. 7, 2024).
Expert Report of M. Weinberg (Jun. 7, 2024).

B. *Deposition Transcripts*

[redacted]
Deposition of J. Giles (Nov. 6, 2020), GOOG-AT-MDL-007172126.
[redacted]
Deposition of N. Jayaram (Sep. 17, 2021), GOOG-AT-MDL-007173084.
[redacted]
Deposition of T. Lipus (Apr. 3, 2024).
[redacted]
Deposition of B. Rowley (Jul. 22, 2021), GOOG-AT-MDL-007177040.
[redacted]
Deposition of S. Spencer (Aug. 12, 2021), GOOG-AT-MDL-007178292.

C. *Produced Google Documents*

1. *Data and Supporting Material*

GOOG-AT-DOJ-DATA-000066771 to GOOG-AT-DOJ-DATA-000066772.
GOOG-AT-EDTX-DATA-000000000 to GOOG-AT-EDTX-DATA-000258388.
GOOG-AT-EDTX-DATA-000276098 to GOOG-AT-EDTX-DATA-001116097.
GOOG-AT-EDTX-DATA-001116098.
GOOG-AT-EDVA-DATA-000000006.
GOOG-AT-MDL-DATA-000486626 to GOOG-AT-MDL-DATA-000488277.
GOOG-AT-MDL-DATA-000561263 to GOOG-AT-MDL-DATA-000561420.
Letter from D. Pearl to B. Nakamura and M. Freeman (Dec. 7, 2023), GOOG-AT-MDL-C-000012885.
Letter from D. Pearl to K. Garcia (Oct. 6, 2023), GOOG-AT-MDL-C-000012826.

Letter from D. Pearl to M. Freeman (Jun. 9, 2023), GOOG-AT-MDL-C-000012751.
Letter from D. Pearl to M. Freeman (Sep. 8, 2023), GOOG-AT-MDL-C-000012795.
Letter from J. Elmer to J. Hogan (Mar. 31, 2022), GOOG-AT-MDL-007334120.
Letter from J. Elmer to J. Hogan (Aug. 19, 2022), GOOG-AT-MDL-007334131.

2. Other Documents

GOOG-AT-MDL-001094067
GOOG-AT-MDL-001397473
GOOG-AT-MDL-002293467
GOOG-AT-MDL-003465605
GOOG-AT-MDL-003995286
GOOG-AT-MDL-004016180
GOOG-AT-MDL-004242638
GOOG-AT-MDL-004555239
GOOG-AT-MDL-006196134
GOOG-AT-MDL-007334120
GOOG-AT-MDL-007374059
GOOG-AT-MDL-007375273
GOOG-AT-MDL-008842383
GOOG-AT-MDL-008842393
GOOG-AT-MDL-008930706
GOOG-AT-MDL-008935836
GOOG-AT-MDL-009293665
GOOG-AT-MDL-009590288
GOOG-AT-MDL-009644018
GOOG-AT-MDL-009644112
GOOG-AT-MDL-009644157
GOOG-AT-MDL-009644182
GOOG-AT-MDL-009644201
GOOG-AT-MDL-009644236
GOOG-AT-MDL-009644238
GOOG-AT-MDL-009644401
GOOG-AT-MDL-009644409
GOOG-AT-MDL-010338120
GOOG-AT-MDL-010836318
GOOG-AT-MDL-011234683
GOOG-AT-MDL-011687180
GOOG-AT-MDL-012701069
GOOG-AT-MDL-015521456
GOOG-AT-MDL-B-001391461
GOOG-AT-MDL-B-003180112
GOOG-AT-MDL-B-004582905
GOOG-AT-MDL-C-000017969
GOOG-AT-MDL-C-000017971
GOOG-AT-MDL-C-000016753
GOOG-AT-MDL-C-000035250

GOOG-AT-MDL-C-000035251
GOOG-AT-MDL-C-000035252
GOOG-AT-MDL-C-000073682
GOOG-DOJ-01805208
GOOG-DOJ-02854344
GOOG-DOJ-03366145
GOOG-DOJ-03366173
GOOG-DOJ-03516570
GOOG-DOJ-03619484
GOOG-DOJ-03643284
GOOG-DOJ-03792729
GOOG-DOJ-03876025
GOOG-DOJ-03998505
GOOG-DOJ-04306227
GOOG-DOJ-04430492
GOOG-DOJ-04602757
GOOG-DOJ-04934481
GOOG-DOJ-04937154
GOOG-DOJ-05270417
GOOG-DOJ-05272070
GOOG-DOJ-05276941
GOOG-DOJ-05282625
GOOG-DOJ-05283173
GOOG-DOJ-05326023
GOOG-DOJ-06525908
GOOG-DOJ-06563186
GOOG-DOJ-06732979
GOOG-DOJ-06818412
GOOG-DOJ-06842351
GOOG-DOJ-06875572
GOOG-DOJ-06882418
GOOG-DOJ-06885161
GOOG-DOJ-07235914
GOOG-DOJ-07825115
GOOG-DOJ-09163089
GOOG-DOJ-09457868
GOOG-DOJ-09459336
GOOG-DOJ-09494195
GOOG-DOJ-09713317
GOOG-DOJ-09714662
GOOG-DOJ-09875881
GOOG-DOJ-09875989
GOOG-DOJ-09916352
GOOG-DOJ-10572595
GOOG-DOJ-10924270
GOOG-DOJ-10924698

GOOG-DOJ-11247631
GOOG-DOJ-11406673
GOOG-DOJ-11733552
GOOG-DOJ-11781854
GOOG-DOJ-11782962
GOOG-DOJ-11790760
GOOG-DOJ-12025827
GOOG-DOJ-12038253
GOOG-DOJ-12059682
GOOG-DOJ-12439154
GOOG-DOJ-12443562
GOOG-DOJ-12700489
GOOG-DOJ-12848608
GOOG-DOJ-12948968
GOOG-DOJ-13199584
GOOG-DOJ-13199603
GOOG-DOJ-13199910
GOOG-DOJ-13199952
GOOG-DOJ-13200158
GOOG-DOJ-13200480
GOOG-DOJ-13201465
GOOG-DOJ-13202659
GOOG-DOJ-13203511
GOOG-DOJ-13204346
GOOG-DOJ-13205869
GOOG-DOJ-13206921
GOOG-DOJ-13207875
GOOG-DOJ-13208074
GOOG-DOJ-13208758
GOOG-DOJ-13208800
GOOG-DOJ-13209957
GOOG-DOJ-13211589
GOOG-DOJ-13212948
GOOG-DOJ-13221355
GOOG-DOJ-13226855
GOOG-DOJ-13227256
GOOG-DOJ-13235100
GOOG-DOJ-13247322
GOOG-DOJ-13364449
GOOG-DOJ-13469175
GOOG-DOJ-13470118
GOOG-DOJ-13550075
GOOG-DOJ-13564564
GOOG-DOJ-13579782
GOOG-DOJ-13605152
GOOG-DOJ-13619370

GOOG-DOJ-13625417
GOOG-DOJ-13627809
GOOG-DOJ-13949282
GOOG-DOJ-13979867
GOOG-DOJ-14000011
GOOG-DOJ-14024199
GOOG-DOJ-14028590
GOOG-DOJ-14029750
GOOG-DOJ-14030931
GOOG-DOJ-14039426
GOOG-DOJ-14298902
GOOG-DOJ-14398809
GOOG-DOJ-14421383
GOOG-DOJ-14514871
GOOG-DOJ-14549757
GOOG-DOJ-14550102
GOOG-DOJ-14566285
GOOG-DOJ-14952787
GOOG-DOJ-15116057
GOOG-DOJ-15119015
GOOG-DOJ-15127000
GOOG-DOJ-15130321
GOOG-DOJ-15208416
GOOG-DOJ-15226550
GOOG-DOJ-15232606
GOOG-DOJ-15247796
GOOG-DOJ-15254730
GOOG-DOJ-15264552
GOOG-DOJ-15389438
GOOG-DOJ-15407321
GOOG-DOJ-15416614
GOOG-DOJ-15417688
GOOG-DOJ-15419945
GOOG-DOJ-15423317
GOOG-DOJ-15423462
GOOG-DOJ-15426012
GOOG-DOJ-15442474
GOOG-DOJ-15588979
GOOG-DOJ-15637938
GOOG-DOJ-15730729
GOOG-DOJ-28385887
GOOG-DOJ-32261273
GOOG-DOJ-AT-00045716
GOOG-DOJ-AT-00070433
GOOG-DOJ-AT-00173317
GOOG-DOJ-AT-00573492

GOOG-DOJ-AT-00608572
GOOG-DOJ-AT-00849635
GOOG-DOJ-AT-01027937
GOOG-DOJ-AT-01133273
GOOG-DOJ-AT-01363996
GOOG-DOJ-AT-01502155
GOOG-DOJ-AT-02199478
GOOG-DOJ-AT-02204351
GOOG-DOJ-AT-02218994
GOOG-DOJ-AT-02425378
GOOG-DOJ-AT-02467209
GOOG-DOJ-AT-02471119
GOOG-DOJ-AT-02471194
GOOG-DOJ-AT-02472888
GOOG-DOJ-AT-02480338
GOOG-DOJ-AT-02512863
GOOG-DOJ-AT-02524665
GOOG-NE-04599495
GOOG-TEX-00000001
GOOG-TEX-00048091
GOOG-TEX-00054839
GOOG-TEX-00090969
GOOG-TEX-00458239
GOOG-TEX-00537828
GOOG-TEX-00682264
GOOG-TEX-00782851
GOOG-TEX-00831373

D. Produced Third-Party Documents

[redacted]

[redacted]

E. Publicly Available Sources

1. Publications

Akbarpour, M., & Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, *88*(2), 425-467.

Ausubel, L. M., & Milgrom, P. R. (2006). The lovely but lonely Vickrey auction. In *Combinatorial Auctions* (P. Cramton, Y. Shoham, & R. Steinberg, Eds.), *MIT Press*.

Balseiro, S. R., & Gur, Y. (2019). Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, *65*(9), 3952-3968.

Balseiro, S R.., Kim, A., Mahdian, M., & Mirrokni, V. (2017). Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*, 15-23.

Barlow, R. E., & Proschan, F. (1975). *Statistical theory of reliability and life testing: probability models* (Vol. 1). New York: Holt, Rinehart and Winston.

Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, *105*(3), 921-57.

Bergemann, D., & Välimäki, J. (2019). Dynamic mechanism design: An introduction. *Journal of Economic Literature*, *57*(2), 235-274.

Cai, Y., Daskalakis, C., & Weinberg, S. M. (2013). Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 618-627.

Choi, A. H. (2009). A Rent Extraction Theory of Right of First Refusal. *The Journal of Industrial Economics*, *57*(2), 252–264.

Choi, H., & Mela, C. F. (2023). Optimizing reserve prices in display advertising auctions. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4523022.

Choi, H., Mela, C. F., Balseiro, S. R., & Leary, A. (2020). Online display advertising markets: A literature review and future directions. *Information Systems Research*, *31*(2), 556-575.

D'Annunzio, A., & Russo, A. (2024). Intermediaries in the online advertising market. *Marketing Science*, *43*(1), 33-53.

Despotakis, S., Ravi, R., & Sayedi, A. (2021). First-price auctions in online display advertising. *Journal of Marketing Research*, *58*(5), 888-907.

Doraszelski, U., Lewis, G., & Pakes, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, *108*(3), 565-615.

Freixas, X., Guesnerie, R., & Tirole, J. (1985). Planning under incomplete information and the ratchet effect. *The Review of Economic Studies*, *52*(2), 173-91.

Fudenberg, D., & Tirole, J. (1991). *Game theory.* MIT Press.

Gans, J.S., & Holden, R. T. (2022). Mechanism design approaches to blockchain consensus, (NBER Working Paper No. w30189). National Bureau of Economic Research.

Goke, S., Weintraub, G. Y., Mastromonaco, R., & Seljan, S. (2022). Bidders' responses to auction format change in internet display advertising auctions. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4021221.

Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first‐price auctions. *Econometrica*, *68*(3), 525-574.

Gui, G., Nair, H., & Niu, F. (2021). Auction throttling and causal inference of online advertising effects. arXiv preprint arXiv:2112.1515.

Harsanyi, J.C. (1967). Games with incomplete information played by Bayesian players, I-III. Part I. The basic model. *Management Science*, *14*(3), 159-182.

Heymann, B. (2020). How to bid in unified second-price auctions when requests are duplicated. *Operations Research Letters*, *48*(4), 446-451.

Hortaçsu, A., Luco, F., Puller, S. L., & Zhu, D. (2019). Does strategic ability affect efficiency? Evidence from electricity markets. *American Economic Review*, *109*(12), 4302-4342.

Hortaçsu, A., & Puller, S. L. (2008). Understanding strategic bidding in multi‐unit auctions: A case study of the Texas electricity spot market. *The RAND Journal of Economics*, *39*(1), 86-114.

Kahan, M., Leshem, S., & Sundaram, R. K. (2012). First-Purchase Rights: Rights of First Refusal and Rights of First Offer. *American Law and Economics Review*, *14*(2), 331–371.

Kaplan, T. R., & Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economitc Theory*, *50*(2), 269-302.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, *13*(3), 227-86.

Kojima, F., & Pathak, P.A. (2009). Incentives and stability in large two-sided matching markets. *American Economic Review*, *99*(3), 608-627.

Krishna, V. (2010). *Auction theory* (2nd ed.). Academic Press.

Levin, J., & Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review: Papers & Proceedings*, *100*(2), 603-07.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory.* Oxford University Press.

McAfee, R. P., & McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, *25*(2), 699-738.

Milgrom, P. R. (2017). *Discovering prices: auction design in markets with complex constraints.* Columbia University Press.

Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press.

Milgrom, P. R., & Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, *128*(1), 1-31.

Milgrom, P. R., and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, *70*(2), 583-601.

Milgrom, P. R., & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157-180.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58-73.

Ostrom, E. (2008). Tragedy of the commons. In *The New Palgrave Dictionary of Economics* (2nd ed.) (S. N. Durlauf & L. E. Blume, Eds.). Palgrave Macmillan.

Ostrovsky, M., & Schwarz, M. (2023). Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy*, *131*(12), 3352-3376.

Pachilakis, M., Papadopoulos, P., Markatos, E. P., & Kourtellis, N. (2019). No more chasing waterfalls: a measurement study of the header bidding ad-ecosystem. In *Proceedings of the Internet Measurement Conference*, 280-293.

Perrigne, I., & Vuong, Q. (2019). Econometrics of auctions and nonlinear pricing. *Annual Review of Economics*, *11*, 27-54.

Reiley, D. H. (2006). Field experiments on the effects of reserve prices in auctions: More magic on the internet. *The RAND Journal of Economics*, *37*(1), 195-211.

Rhuggenaath, J., Akcay, A., Zhang, Y., & Kaymak, U. (2022). Setting reserve prices in second-price auctions with unobserved bids. *INFORMS Journal on Computing*, *34*(6), 2950-2967.

Riley, J. G., & Samuelson, W. F. (1981). Optimal auctions. *American Economic Review*, *71*(3), 381-392.

Roth, A.E. (2015). *Who gets what—and why: The new economics of matchmaking and market design.* Houghton Mifflin Harcourt.

Rothkopf, M.H., Teisberg, T.J., & Kahn, E.P. (1990). Why are Vickrey auctions rare? *Journal of Political Economy*, *98*(1), 94-109.

Scott, D. W. (1992). *Multivariate density estimation: Theory, practice, and visualization.* John Wiley & Sons.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37.

Xu, J., Lee, K. C., Li, W., Qi, H., & Lu, Q. (2015). Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2217-2226.

Ye, Z., Zhang, D.J., Zhang, H., Zhang, R., Chen, X., & Xu, Z. (2023). Cold start to improve market thickness on online advertising platforms: Data-driven algorithms and field experiments. *Management Science*, *69*(7), 3838-60.

Zhang, W., Kitts, B., Han, Y., Zhou, Z., Mao, T., He, H., Pan, S., Flores, A., Gultekin, S., & Weissman, T. (2021). Meow: A space-efficient nonparametric bid shading algorithm. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3928-3936.

2. *Google News Articles, Blog, and Help Pages*

Bellack, J., Exchange Bidding now available to all customers using DoubleClick for Publishers, Google Ad Manager (Apr. 4, 2018), https://blog.google/products/admanager/exchange-bidding-now-available-to-a.

Bellack, J., Smarter optimizations to support a healthier programmatic market, Google Ad Manager (May 12, 2016), https://blog.google/products/admanager/smarter-optimizations-to-suppor.

Bigler, J., An update on first price auctions for Google Ad Manager, Google Ad Manager (May 10, 2019), https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager.

Bigler, J. Rolling out first price auctions to Google Ad Manager partners, Google Ad Manager (Sep. 5, 2019), https://blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners.

Cox, S., Simplifying programmatic: first price auctions for Google Ad Manager, Google Ad Manager (Mar. 6, 2019), https://blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager.

Google, 10 Tips for Google Ads Budget Management, Google Ads Resources (Mar. 20, 2023), https://ads.google.com/intl/en_us/home/resources/articles/stretching-your-google-ads-budget.

Google, 2019 Google Ad Manager releases archive, Google Ad Manager Help (Jul. 1, 2019), https://support.google.com/admanager/answer/9197913?hl=en#expand&expand070119.

Google, 2022 Google Ad Manager releases archive, Google Ad Manager Help, https://support.google.com/admanager/answer/11586212?hl=en#zippy=%2Cjune-optimize-pricing-video-protections-troubleshooting-for-mcm-google-analytics-integration-webview-api-for-ads-updates-to-bid-rejection-reason.

Google, Ad Manager report metrics, Google Ad Manager Help (accessed Sep. 29, 2023), https://support.google.com/admanager/table/7568664?hl=en.

Google, Ad-serving Protections for Teens, Google Ad Manager and Ad Exchange Program Policies (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/12171027?hl=en&sjid=2354124647482356745-EU.

Google, Bid data sharing, Authorized Buyers Help (accessed Jul. 24, 2024), https://support.google.com/authorizedbuyers/answer/2696468?hl=en.

Google, Bids data in Ad Manager Data Transfer (Beta), Google Ad Manager Help (accessed Jul. 24, 2024), https://support.google.com/admanager/answer/7357436?hl=en.

Google, Block general categories, Google Ad Manager Help (accessed Jul. 17, 2024), https://support.google.com/admanager/answer/2913554?sjid=261789886268128879-EU.

Google, Block sensitive categories, Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/2541069?sjid=14973969736947852707-EU.

Google, Coming soon: April 17, 2023 edition, Display & Video 360 announcements (Apr. 17, 2023), https://support.google.com/displayvideo/answer/13511229.

Google, Create a campaign, Display & Video 360 Help (accessed Mar. 12, 2024), https://support.google.com/displayvideo/answer/7205081?hl=en.

Google, Create a Display campaign, Google Ads Help (accessed Mar. 12, 2024), https://support.google.com/google-ads/answer/10759203.

Google, Custom bidding overview and limitations, Display & Video 360 Help (accessed Oct. 31, 2023), https://support.google.com/displayvideo/answer/9723477?hl=en.

Google, Delivery basics: Ad competition with dynamic allocation, Google Ad Manager Help (accessed Jan. 8, 2024), https://support.google.com/admanager/answer/3721872?hl=en.

Google, Determine a bid strategy based on your goals, Google Ads Help (accessed Oct. 31, 2023), https://support.google.com/google-ads/answer/2472725.

Google, DFP and dynamic allocation, DoubleClick for Publishers Help (captured on Sep. 22, 2015), https://web.archive.org/web/20150922150140/https://support.google.com/dfp_premium/answer/3447903.

Google, Display & Video 360: An integrated solution for end-to-end advertising campaigns, Display & Video 360 (accessed Oct. 31, 2023), https://services.google.com/fh/files/misc/display_and_video_360_product_overview.

Google, Display & Video 360 Overview (accessed Jan. 21, 2024), https://support.google.com/displayvideo/answer/9059464?hl=en.

Google, Enhanced Automation, Display & Video 360 Help (accessed Oct. 31, 2023), https://support.google.com/displayvideo/answer/6130826?hl=en.

Google, Evaluate and Optimize Your Bids, Google Ads Help (accessed Oct. 31, 2023), https://support.google.com/google-ads/answer/7085711?hl=en.

Google for Developers, Cookie Matching, Real-time Bidding (accessed Jul. 17, 2024), https://developers.google.com/authorized-buyers/rtb/cookie-guide.

Google, Get Comprehensive Yield Management with Google Ad Manager, Google Ad Manager (accessed Oct. 25, 2023), https://admanager.google.com/home/resources/feature-brief-yield-management.

Google, Google Ads (accessed Jan. 5, 2024), https://ads.google.com/intl/en_us/home.

Google, Google Closes Acquisition of DoubleClick, News from Google (Mar. 11, 2008), http://googlepress.blogspot.com/2008/03/google-closes-acquisition-of_11.

Google, Google Platform Services Terms and Conditions, Google (accessed Sep. 27, 2023), https://www.google.com/google-ad-manager/platform/terms.

Google, How Active View metrics are calculated, Google Ad Manager Help (accessed Sept. 29, 2023), https://support.google.com/admanager/answer/6233478?hl=en&ref_topic=7506202.

Google, How Authorized Buyers Work With Google Ad Manager, Google Ad Manager Resources (accessed Sep. 27, 2023), https://admanager.google.com/home/resources/how_authorized_buyers_work_with_google.

Google, Impressions: Definition, Google Ads Help (accessed Dec. 16, 2023), https://support.google.com/google-ads/answer/6320?hl=en.

Google, Improved header bidding support in Google Ad Manager, Google Ad Manager Blog (Apr. 27, 2022), https://blog.google/products/admanager/improved-header-bidding-support-in-google-ad-manager.

Google, Introduction to Open Bidding, Google Ad Manager Help (accessed Oct. 31, 2023), https://support.google.com/admanager/answer/7128453?hl=en.

Google, Line item types and priorities, Google Ad Manager Help (accessed Jan. 7, 2024), https://support.google.com/admanager/answer/177279?hl=en.

Google, Optimize floor prices in unified pricing rules (Beta), Google Ad Manager Help, https://support.google.com/admanager/answer/11385824.

Google, Run a Manual Experiment, Google Ad Manager Help (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/9799933?hl=en#zippy=%2Cuser-messages%2Cnative-ad-style%2Cunified-pricing-rules%2Cyield-groups%2Cheader-bidding-trafficking-for-prebid.

Google, Set a fixed CPM bid for a line item, Display & Video 360 Help (accessed Jul. 24, 2024), https://support.google.com/displayvideo/answer/2696858?hl=en.

Google, System maximums and limits, Google Ad Manager Help Center (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/1628457?hl=en.

Google, Tag an ad request for child-directed treatment (TFCD), Google Ad Manager and Ad Exchange Program Policies (accessed Jun. 24, 2024), https://support.google.com/admanager/answer/3671211?hl=en&ref_topic=28145.

Google, Targeting types, Google Ad Manager Help (accessed Jun. 26, 2024), https://support.google.com/admanager/answer/2884033?sjid=5613954590018599404-NC#custom-targeting.

Google, Test Campaigns with Ease with Ads Experiments, Google Ads (accessed Jun. 26, 2024), https://ads.google.com/intl/en_us/home/tools/experiments.

Google, [UA] Latency and why it impacts Google Ads Clicks and Analytics Sessions, Google Analytics Help (accessed Jul. 24, 2024), https://support.google.com/analytics/answer/4589209?hl=en.

Google, Understand Video Solutions report metrics, Google Ad Manager Help (accessed Jul. 24, 2024), https://support.google.com/admanager/answer/2759433#video-viewership.

Google, Unified First-Price Auction - Best practices, Google Ad Manager (accessed Jun. 15, 2024), https://services.google.com/fh/files/misc/unified_first-price_auction_best_practices.

Google, Unified pricing rules, Google Ad Manager Help (accessed Jun. 25, 2024), https://support.google.com/admanager/answer/9298008.

Google, Use frequency capping, Google Ads Help (accessed Jul. 24, 2024), https://support.google.com/google-ads/answer/6034106.

Google, Value CPM, Google Ad Manager (accessed Oct. 15, 2023), https://support.google.com/admanager/answer/177222?hl=en.

Hedrick, L., Get Actionable Measurement with Display & Video 360's Insights Module (Nov. 15, 2018), Google Marketing Platform, https://blog.google/products/marketingplatform/360/get-actionable-measurement-display-video-360s-insights-module.

Hsiao, S., How our display buying platforms share revenue with publishers, Google Ad Manager (Jun. 23, 2020), https://blog.google/products/admanager/display-buying-share-revenue-publishers.

Levitte, G., Improved header bidding support in Google Ad Manager, Google Ad Manager (Apr. 27, 2022), https://blog.google/products/admanager/improved-header-bidding-support-in-google-ad-manager.

Ramaswamy, S., Introducing simpler brands and solutions for advertisers and publishers, Google Blog (Jun. 27, 2018), https://blog.google/technology/ads/new-advertising-brands.

Sivasubramanian, V., Help me with Waterfall setup and understanding Passback Tags, Google Ad Manager Help (Mar. 5, 2019), https://support.google.com/admanager/thread/2065360/help-me-with-waterfall-setup-and-understanding-pa ssback-tags?hl=en.

3. News Articles

AdExchanger, Google's Scott Spencer On DoubleClick Ad Exchange Auction And Data Management, AdExchanger (Feb. 9, 2010), https://www.adexchanger.com/ad-exchange-news/googles-scott-spencer-on-doubleclick-ad-exchange-auction-and-data-management.

Benes, R. Ad buyer, beware: How DSPs sometimes play fast and loose, Digiday (May 25, 2017), https://digiday.com/marketing/dsp-squeeze-buyers.

Benes, R., Five Charts: The State of Header Bidding, eMarketer Insider Intelligence (May 30, 2019), https://www.insiderintelligence.com/content/five-charts-the-state-of-header-bidding.

Benes, R., Unraveling header bidding's problems with user data, Digiday (Mar. 20, 2017), https://digiday.com/media/header-bidding-security.

Chen, Y., WTF is supply path optimization?, Digiday (May 22, 2023), https://digiday.com/media/what-is-supply-path-optimization.

Cross, T., Who Invented What in Ad Tech? – Part Two, VideoWeek (Jun. 26, 2018), https://videoweek.com/2018/06/26/who-invented-what-in-ad-tech-part-two.

Curran, J., Opinion: For Publishers, Header Bidding Discrepancies Can Outweigh Revenue Lift, AdExchanger (Jul. 8, 2016), https://www.adexchanger.com/the-sell-sider/publishers-header-bidding-discrepancies-can-outweigh-revenue-lift.

Davies, J., What to know about Google's implementation of first-price ad auctions, Digiday (Sep. 6, 2019), https://digiday.com/media/buyers-welcome-auction-standardization-as-google-finally-goes-all-in-on-first-price.

Digiday, In programmatic, buyers sometimes don't know what type of auction they're bidding in, Digiday (Jun. 30, 2017), https://digiday.com/marketing/ad-buyers-programmatic-auction.

Jatain, V., Understanding Header Bidding And How To Leverage It, Forbes (Sep. 17, 2019), https://www.forbes.com/sites/forbescommunicationscouncil/2019/09/17/understanding-header-bidding-and-how-to-leverage-it/?sh=332097315c18.

Joseph, S., Ad-buying agencies are cozying up to SSPs, creating more transparency questions, Digiday (Mar. 31, 2020), https://digiday.com/marketing/ad-buying-agencies-are-cozying-up-to-ssps-creating-more-transparency-questions.

Joseph, S., DSPs are cracking down on bid duplication, Digiday (May 12, 2020), https://digiday.com/media/dsps-are-cracking-down-on-bid-duplication.

Joseph, S., To reduce auction duplication, buyers start to enforce sellers.json, Digiday (Oct. 16, 2019), https://digiday.com/media/reduce-auction-duplication-buyers-start-enforce-sellers-json.

Joseph, S., WTF are post-auction discounts?, Digiday (Apr. 2 2020), https://digiday.com/media/post-auction-discounts.

Kaplan, D., On Ad Networks: Pork Bellies, Diamonds, Or The New Direct Marketing?, Forbes (Apr. 8, 2008), https://www.forbes.com/2008/04/08/online-ad-networks-tech-cx_pco_0408paidcontent.html?sh=7414ef02cb8e.

Ong, J., Adtech firm OpenX unveils an industry-changing fusion of real-time bidding and ad networks, The Next Web News (Jun. 9, 2014), https://thenextweb.com/news/adtech-firm-openx-unveils-industry-changing-fusion-real-time-bidding-ad-networks.

Poncz, O., Traffic duplication might be a bigger problem than ad fraud, AdExchanger (Jan. 11, 2016), https://www.adexchanger.com/data-driven-thinking/traffic-duplication-might-even-be-a-bigger-problem-than-ad-fraud.

Sluis, S., Attack Of The Clones: Programmatic's Hidden Scourge Of Bid Duplication, AdExchanger (Jan. 17, 2024), https://www.adexchanger.com/platforms/attack-of-the-clones-programmatics-hidden-scourge-of-bid-duplication.

Sluis, S., Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price, AdExchanger (Sep. 5, 2017), https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price.

Sluis, S., Everything You Need To Know About Bid Shading, AdExchanger (Mar. 15, 2019), https://www.adexchanger.com/online-advertising/everything-you-need-to-know-about-bid-shading.

Sluis, S., Explainer: More On The Widespread Fee Practice Behind The Guardian's Lawsuit Vs. Rubicon Project, AdExchanger (Mar. 30, 2017),

https://www.adexchanger.com/ad-exchange-news/explainer-widespread-fee-practice-behind-guardians-lawsuit-vs-rubicon-project.

Sluis, S., Google Ad Manager Builds A Bridge To Prebid—But Don't Call It A Two-Way Street, AdExchanger (Apr. 27, 2022), https://www.adexchanger.com/platforms/google-ad-manager-builds-a-bridge-to-prebid-but-dont-call-it-a-two-way-street.

Sluis, S., Google Ad Manager Policy Changes Don't Hurt Publishers, According to Advertiser Perceptions, AdExchanger (May 5, 2020), https://www.adexchanger.com/platforms/google-ad-manager-policy-changes-dont-hurt-publishers-according-to-advertiser-perceptions.

Sluis, S., Google Removes Its 'Last-Look' Auction Advantage, AdExchanger (Mar. 31, 2017), https://www.adexchanger.com/platforms/google-removes-last-look-auction-advantage.

Sluis, S., Header Bidding Unleashed a Huge Infrastructure Problem and Ad Tech Will Either Sink or Swim, AdExchanger (Apr. 24, 2017), https://www.adexchanger.com/platforms/header-bidding-unleashed-huge-infrastructure-problem-ad-tech-will-either-sink-swim.

Sluis, S., OpenX Lays Off 100 Employees And Pivots To Video, AdExchanger (Dec. 18, 2018), https://www.adexchanger.com/platforms/openx-lays-off-100-employees-and-pivots-to-video.

Sluis, S., The Trade Desk suppresses bid duplication amid COVID-19 traffic surge, AdExchanger (Apr. 21, 2020), https://www.adexchanger.com/platforms/the-trade-desk-suppresses-bid-duplication-amid-covid-19-traffic-surge.

Story, L., DoubleClick to Set Up an Exchange for Buying and Selling Digital Ads, New York Times (Apr. 4, 2007), https://www.nytimes.com/2007/04/04/business/media/04adco.html.

4. *Other Websites and Material*

Abhilasha, The Ultimate Guide to Open Bidding for Publishers, headerbidding.co (Jul. 25, 2023), https://headerbidding.co/open-bidding-ultimate-guide.

Ad Ops Insider, Header Bidding Explained Step-by-Step (Jun. 8, 2015), https://www.adopsinsider.com/header-bidding/header-bidding-step-by-step.

Ad.Plus, Header Bidding vs Open Bidding, Ad.Plus Blog (Jan. 6, 2023), https://blog.ad.plus/header-bidding-vs-open-bidding.

AdPushup, Header Bidding (2023), https://www.adpushup.com/header-bidding-guide.

AffiliateSeeking.com, Ad Networks (captured on Jan. 16, 2008), https://web.archive.org/web/20080116101025/https://www.affiliateseeking.com/list/23000001/1.

Amazon, Unified Ad Marketplace, Amazon Publisher Services (accessed Sep. 12, 2023), https://aps.amazon.com/aps/unified-ad-marketplace/index.

American Economic Association, Paul Milgrom, Distinguished Fellow 2020, American Economic Association (2020), https://www.aeaweb.org/about-aea/honors-awards/distinguished-fellows/paul-milgrom.

Athey, S., Economics 980: Topics in Market Design, Stanford University (accessed Sept. 27, 2023), https://gsb-faculty.stanford.edu/susan-athey/economics-980-topics-market-design.

Auctionomics and Power Auctions, Incentive Auction Rules Option and Discussion, FCC (Sep. 12, 2012), https://docs.fcc.gov/public/attachments/FCC-12-118A2.

Bentahar, A., Bid Adjustments Simplified: Run Fair Auctions with no Hassle, Pubstack (May 2, 2021), https://www.pubstack.io/topics/bid-adjustments-simplified.

Bhatia, S., Google Open Bidding (OB) Explained, AdSparc (Apr. 12, 2022), https://adsparc.com/google-open-bidding-explained.

Bidgear, A Well Organized Passback Strategy (Aug. 7, 2016), https://bidgear.com/blog/a-well-organized-passback-strategy-24.

Clickio, Monetization (accessed Jul. 7, 2024), https://clickio.com/monetization.

Criteo S.A., Form 10-K, United States Securities and Exchange Commission (2022), https://www.sec.gov/Archives/edgar/data/1576427/000157642723000023/crto-20221231.

Criteo, What You Need to Know About First-Price Auctions and Criteo, Criteo Updates (Sep. 26, 2023), https://www.criteo.com/blog/first-price-auctions.

Das, H., Header Bidding vs Open Bidding – What Should You Know!, headerbidding.co Blog (Feb. 23, 2024), https://headerbidding.co/header-bidding-vs-open-bidding.

Doherty, W., The door in the floors: Transparency, price discovery, and market efficiency in programmatic, The Current (Nov. 8, 2023), https://www.thecurrent.com/will-doherty-transparency-programmatic-data.

DoubleClick for Publishers, Terminology differences with DART, DoubleClick for Publishers Help (captured Feb. 11, 2012), https://web.archive.org/web/20120211164005/http://support.google.com/dfp_premium/bin/answer.py?hl=n&answer=158833.

DoubleVerify, Latency in Digital Advertising: A Guide for Publishers, DV Publisher Insights (Oct. 21, 2019), https://pub.doubleverify.com/blog/latency-in-digital-advertising-a-guide-for-publishers.

INFORMS, About AMD - Auctions and Market Design (accessed Sept. 27, 2023), http://connect.informs.org/auctionsandmarketdesign/about-us/aboutamd.

IAB Tech Lab, OpenRTB (Jul. 27, 2020), https://iabtechlab.com/standards/openrtb.

Interactive Advertising Bureau, Glossary: Digital Media Buying & Planning (Apr. 2016), https://www.iab.com/wp-content/uploads/2016/04/Glossary-Formatted.

Laurendon, L., Google Unified First Price Auction Explained, Smart Ad Server Blog (captured on Mar. 7, 2022), http://web.archive.org/web/20220307203150/https://smartadserver.com/articles/google-unified-first-price-auction-explained.

Magnite Team, Magnite Unveils New Demand Manager Feature Powered by Machine Learning to Help Publishers Earn Incremental Revenue, Magnite (Oct. 5, 2023), https://www.magnite.com/press/magnite-unveils-demand-manager-machine-learning-feature.

Media Rating Council, Minimum Standards for Media Rating Research, Media Rating Council (accessed Sep. 29, 2023), https://mediaratingcouncil.org/standards-and-guidelines.

Meta, Code of Conduct, Meta Audience Network (accessed Dec. 2, 2023), https://www.facebook.com/audiencenetwork/partner-program/code-of-conduct.

Microsoft, Discover the possibilities with experiments, Microsoft Advertising (accessed Jun. 26, 2024), https://help.ads.microsoft.com/apex/index/3/en/56908.

Milgrom, P. R., Cunningham, S. J., & Beck, M. R. (2023). U.S. Patent No. 11,574,358-B2. Washington, DC: U.S. Patent and Trademark Office.

Munson, B., Comcast's FreeWheel launches unified ad decisioning, StreamTV Insider (Apr. 15, 2020), https://www.streamtvinsider.com/tech/comcast-s-freewheel-launches-unified-ad-decisioning.

Naylon, Jenna, Media Planning and Buying Strategies: The Importance of Test & Learn, MatrixPoint (accessed Jul. 17, 2024), https://www.thematrixpoint.com/resources/articles/media-planning-and-buying-strategies-the-importance-of-test-learn.

NBER, Market Design (2023), https://www.nber.org/programs-projects/programs-working-groups%23Groups/market-design.

Nolet, M., RTB Part II: Supply supply supply!, Mike On Ads (Sep. 19, 2009), http://www.mikeonads.com/2009/09/19/rtb-part-ii-supply-supply-supply.

OpenX, Google & OpenX Release Study Showing Publisher Partners Experience 48% Revenue Lift Through Google Exchange Bidding Collaboration (Feb. 15, 2018), https://www.openx.com/press-releases/google-openx-revenue-lift.

OpenX, Selling rules (Dec. 12, 2016), https://web.archive.org/web/20170708051049/https://docs.openx.com/Content/publishers/userguide_inventory_realtimeselling.

OpenX, OpenX transforms concept of SSP with industry-first demand fusion technology (Jun. 9, 2014), https://www.openx.com/press-releases/openx-transforms-concept-of-ssp-with-industry-first-demand-fusion-technology.

Prebid.org, About (accessed Jul. 25, 2024), https://prebid.org/about.

Prebid.org, Prebid.js Ad Ops (accessed Jul. 24, 2024), https://docs.prebid.org/adops/price-granularity.html#prebid-default-price-granularities.

Prebid.org, Prebid.js FAQ (accessed Dec. 1, 2023), https://docs.prebid.org/dev-docs/faq.

Prebid.org, Prebid.js Publisher API Reference (accessed Jul. 24, 2024), https://docs.prebid.org/dev-docs/publisher-api-reference/setConfig.html#setConfig-Cpm-Rounding.

Prebid.org, Price Granularity (accessed Jul. 12, 2024), https://docs.prebid.org/adops/price-granularity.

Prebid.org, Running Prebid.js without an ad server (accessed Sep. 7, 2023), https://docs.prebid.org/dev-docs/examples/no-adserver.

Pubmatic, Pubmatic SSP: Maximize Advertising Revenue and Control How Your Audiences are Accessed, Pubmatic (accessed Oct. 25, 2023), https://pubmatic.com/products/pubmatic-ssp-for-publishers.

Refinitiv Streetevents, Q2 2021 Magnite Inc Earnings Call (Aug. 5, 2021), https://investor.magnite.com/static-files/950ded3e-8953-4cf9-a043-d0e72b1fb856.

Scipy, scipy.spatial.ConvexHull, SciPy v1.11.4 Manual (accessed Jan. 10, 2024), https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.ConvexHull.

SciPy, scipy.stats.gaussian_kde, SciPy v1.11.4 Manual (accessed Dec. 3, 2023), https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.

Semret, N., Introducing Minimum CPM Recommendations on DoubleClick Ad Exchange, DoubleClick Publisher Blog (Nov. 1, 2011), https://doubleclick-publishers.googleblog.com/2011/11/introducing-minimum-cpm-recommendations.

SLMath, Mathematics and Computer Science of Market and Mechanism Design (2023), https://www.slmath.org/programs/333.

Stack Overflow, What are advertising passback tags and common implementation (Feb. 19, 2014), https://stackoverflow.com/questions/19112923/what-are-advertising-passback-tags-and-common-implementation.

Stanford University, Stanford Bulletin (2023), https://explorecourses.stanford.edu/search;jsessionid=zg9iqiunv63g16z0qr48bmb90?q=ECON+136%3a+Market+Design&view=catalog&filter-coursestatus-Active=on&academicYear=20232024.

The Royal Swedish Academy of Sciences, The quest for the perfect auction, Nobelprize.org (2020), https://www.nobelprize.org/uploads/2020/09/popular-economicsciencesprize2020.

The Trade Desk, DFP Features: Integrated planning and activation (accessed Oct. 31, 2023), https://www.thetradedesk.com/us/our-platform/dsp-demand-side-platform/plan-campaigns.

The Trade Desk, Conversion Lift: What it is, how it works, and best practices (Aug. 23, 2023), https://www.thetradedesk.com/us/resource-desk/best-practices-for-better-conversion-lift.

Xandr, Publisher Platforms, Microsoft Advertising (accessed Oct. 25, 2023), https://about.ads.microsoft.com/en-us/solutions/xandr/publisher-platforms-scaled-buying-selling-solutions.

Xandr, Seller Best Practices (May 5, 2023), https://docs.xandr.com/bundle/industry-reference/page/seller-best-practices.

Xandr, Understanding Optimization, Xandr Documentation Center (Feb. 7, 2022), https://docs.xandr.com/bundle/monetize_monetize-standard/page/topics/understanding-optimization.

Yahoo, Verizon Media Is Now Yahoo (Sep. 1, 2021), https://www.advertising.yahooinc.com/post/verizon-media-is-now-yahoo.

Zawadziński, M., How Does Real-Time Bidding (RTB) Work?, Clearcode Blog (Jul. 2, 2021), https://clearcode.cc/blog/real-time-bidding.

Zawadziński, M. & Sweeney, M., What is cookie syncing and how does it work?, Clearcode Blog (Jan. 31, 2024), https://clearcode.cc/blog/cookie-syncing.

Zawadziński, M. & Sweeney, M., What is Waterfalling and How Does it Work?, Clearcode Blog (Sep. 1, 2016), https://clearcode.cc/blog/what-is-waterfalling.